

QM-1
HARDWARE LEVEL
USER'S MANUAL



NANODATA CORPORATION
2457 Wehrle Drive Williamsville, New York 14221

	Q	Q	Q	M		M			1
	Q		Q	MM		MM			11
Q			Q	M	M	M	M		1 1
Q			Q	M		M			1
Q			Q	M		M		-----	1
Q			Q	M		M			1
Q		Q	Q	M		M			1
	Q		Q	M		M			1
	Q	Q	Q	M		M			11111

HARDWARE LEVEL USER'S MANUAL

SECOND EDITION

MARCH 21, 1973

NANODATA CORPORATION
2457 NEHRLE DRIVE
WILLIAMSVILLE, NEW YORK 14221
716-631-5880

COPYRIGHT C 1972 NANODATA CORPORATION

TABLE OF CONTENTS

1	SCOPE AND PURPOSE -----	9
2	INTRODUCTION TO MICROPROGRAMMING AND THE QM-1 -----	10
2.1	BASIC INSTRUCTION SEQUENCE -----	11
2.2	MICROPROGRAMMED CONTROL -----	12
2.3	USES OF MICROPROGRAMMING -----	14
2.4	HORIZONTAL AND VERTICAL CONTROL -----	16
2.5	THE QM-1 CONTROL HIERARCHY -----	18
3	USER AND MACHINE HIERARCHIES -----	20
4	QM-1 FUNCTIONAL SPECIFICATIONS, PART I -----	23
4.1	GENERAL -----	23
4.2	MAJOR RESOURCES AND ORGANIZATION -----	24
4.2.1	MAJOR BUSING STRUCTURE AND LOCAL STORE -----	24
4.2.2	LOCAL STORE SPECIAL FEATURES -----	27
4.2.2.1	MICRO INSTRUCTION REGISTER -----	27
4.2.2.2	MICRO PROGRAM COUNTERS -----	28
4.2.2.3	INDEX ALU FEATURE -----	28
4.2.3	ALU AND SHIFTER -----	29
4.2.3.1	ALU -----	29
4.2.3.2	SINGLE SHIFTS -----	30
4.2.3.3	DOUBLE SHIFTS -----	31
4.2.3.4	CARRY CONTROL -----	32
4.2.3.5	TEST CONDITIONS -----	34
4.2.3.6	SIXTEEN-BIT MODE -----	35
4.2.4	CONTROL STORE -----	36
4.2.5	EXTERNAL STORE -----	37
4.2.6	MAIN STORE -----	39
4.2.6.1	GENERAL -----	39
4.2.6.2	MS OPERATIONS -----	40
4.2.6.3	MS ADDRESSING AND PROTECTION (QM-1 OPTION) -----	42
4.2.6.4	RMI UNIT (QM-1 OPTION) -----	43
4.3	SIX-BIT CONTROL STRUCTURE -----	44
4.3.1	GENERAL -----	44
4.3.2	F-STORE -----	45
4.3.2.1	GENERAL -----	45
4.3.2.2	BUS CONTROL F'S -----	46
4.3.2.3	SPECIAL F'S -----	47

4.3.2.4	G'S	49
4.3.3	AUX FIELDS	50
4.3.4	ALUF (QM-1 OPTION)	52
4.4	TIMING	53
4.5	NANOPROGRAM CONTROL	56
4.5.1	CONTROL MATRIX	56
4.5.2	NANOSTORE ADDRESSING	59
4.5.2.1	PRIORITY SELECT	59
4.5.2.2	PROGRAM CHECK	61
4.5.2.3	NANOBRANCH	62
4.5.2.4	EXTERNAL INTERRUPTS	63
4.5.3	NANOPROGRAM COUNTER	65
4.5.3.1	GENERAL	65
4.5.3.2	MICROINSTRUCTION EXECUTION	66
4.5.4	NANOPROGRAM FLOW	68
4.5.4.1	NANOPROGRAM CONNECTION	68
4.5.4.2	INTERRUPTIBILITY	69
4.5.4.3	HOLD	70
4.6	EXTERNAL INTERFACE	71
4.7	WRITING NANOSTORE	73
4.8	READ-ONLY MEMORIES AND MACHINE START	74
5	QM-1 FUNCTIONAL SPECIFICATIONS, PART II	75
5.1	GENERAL	75
5.2	SUMMARY OF NANOPRIMITIVE CONTROLS	76
5.2.1	K-VECTOR CONTROL FIELDS	76
5.2.2	T-VECTOR CONTROL FIELDS	78
5.3	FUNDAMENTAL TIMING CONSIDERATIONS	82
5.3.1	GENERAL	82
5.3.2	LEADING EDGE FUNCTIONS	84
5.3.3	TRAILING EDGE EVENTS	85
5.3.3.1	T-CLOCK EVENTS	85
5.3.3.2	F-CLOCK EVENTS	85
5.3.3.3	R-CLOCK EVENTS	85
5.3.4	R31 OPERATIONS	86
5.4	MEMORY REFERENCE FUNCTIONS	88
5.4.1	NANOSTORE OPERATIONS	88
5.4.1.1	READ NS	88
5.4.1.2	WRITE NS	89
5.4.2	CONTROL STORE OPERATIONS	90
5.4.2.1	CONTROL STORE ADDRESS SELECTION	90
5.4.2.2	READ CS	91

5.4.2.3	WRITE CS	91
5.4.3	MAIN STORE OPERATIONS	92
5.4.3.1	READ MAIN STORE	92
5.4.3.2	WRITE MAIN STORE	93
5.4.3.3	READ-MODIFY-WRITE	93
5.4.4	EXTERNAL OPERATIONS	94
5.5	DATA TRANSFER FUNCTIONS	95
5.5.1	T TRANSFERS	95
5.5.2	F TRANSFERS	96
5.5.3	R TRANSFERS	99
5.6	DATA MANIPULATION FUNCTIONS	100
5.6.1	GENERAL	100
5.6.2	OPERATION OF THE ARITHMETIC-LOGIC UNIT	101
5.6.3	OPERATION OF THE SHIFTER AND SHIFTER EXTENSION	105
5.6.4	OPERATION OF THE INDEX ALU	107
5.6.5	MPC OPERATIONS	110
5.6.6	INCREMENT F AND DECREMENT F OPERATIONS	111
5.6.7	ALUF OPERATION	112
5.6.8	RMI OPERATION	114
5.7	CONDITIONAL FUNCTIONS	115
5.7.1	TEST FIELDS AND MASKS	115
5.7.2	TESTS	117
5.8	MISCELLANEOUS FUNCTIONS	119
5.8.1	GENERATE INTERRUPT	119
5.8.2	AUXILLARY ACTION	120
5.8.3	MISCELLANEOUS CONTROLS	121
5.9	REFERENCE LAYOUT AND MAPS	122
5.9.1	CPU REGISTER ASSIGNMENT AND LAYOUT	122
5.9.2	NANOSTORE MAP	124
6	NANOPROGRAMMING LANGUAGE SPECIFICATIONS	125
6.1	GENERAL	125
6.2	ELEMENTS	126
6.2.1	SOURCE STATEMENTS	126
6.2.2	CHARACTER SET	126
6.2.3	SYMBOLIC NAMES	127
6.2.4	DELIMITERS	127
6.2.4.1	FIELD DELIMITERS	127
6.2.4.2	SYMBOLIC NAME DELIMITERS	128
6.2.5	COMMENTS	128
6.2.5.1	COMMENT STATEMENTS	128
6.2.5.2	COMMENT FIELDS	128

6.2.6	BLANKS -----	128
6.3	NANOWORD DEFINITION -----	129
6.3.1	LABELS STATEMENTS -----	129
6.3.2	COMMAND STATEMENTS -----	130
6.3.2.1	PSEUDO COMMAND OPERATORS -----	130
6.3.2.2	NANOPRIMITIVE COMMANDS -----	133
6.3.2.3	VARIABLES SYMBOLS AND CONSTANTS -----	141
6.3.2.4	6 BIT DATA TRANSFER COMMANDS -----	143
6.3.3	CONTROL STATEMENTS -----	145
6.3.3.1	ASSEMBLY CONTROL STATEMENTS -----	145
6.3.3.2	LISTING CONTROL STATEMENTS -----	146
6.4	OPERATION (VERSION 1, LEVEL 2, UNDER N. C. S.) -----	147
6.4.1	INVOCATION -----	147
6.4.2	ERROR FLAGS -----	147
7	NANOPROGRAMMING EXAMPLES -----	150
7.1	BASIC MODEL NANOPROGRAMS -----	150
7.2	USE OF PREDEFINED NANOWORDS -----	154
7.3	CONTROL STORE ACCESS NANOPROGRAMS -----	156
7.4	MICROINSTRUCTION BRANCH NANOPROGRAMS -----	159
7.5	ARITHMETIC NANOPROGRAMS (SHIFTS) -----	162
7.6	MULTIWORD NANOPROGRAMS -----	168
8	QM-1 I/O SYSTEM -----	173
8.1	GENERAL -----	173
8.2	QM-1 I/O CONTROLS -----	176
8.2.1	GENERAL -----	176
8.2.2	CPU PORT OPERATION -----	177
8.2.3	CPU INTERRUPT OPERATION -----	180
8.3	STANDARD CHANNEL CONTROLLER -----	181
8.3.1	GENERAL -----	181
8.3.2	CHANNEL COMMANDS -----	187
8.4	STANDARD DEVICE CONTROLLER -----	189
8.4.1	GENERAL -----	189
8.4.2	DEVICE CONTROL -----	190
8.4.2.1	DEVICE CONTROL WORD -----	190
8.4.2.2	STATUS -----	192
8.4.2.3	INTERRUPT MECHANISM -----	196
8.4.2.4	DIRECT REQUEST MECHANISM -----	198
8.4.2.5	DEVICE COMMANDS -----	199
8.4.3	DATA ROUTING SUPPORT -----	201
8.4.3.1	BUFFER ADDRESS REGISTER -----	201

8.4.3.2	WORD COUNT REGISTER -----	202
8.4.3.3	DATA TRANSFER -----	203
8.4.3.4	DATA CHAINING -----	204
8.4.3.5	DIRECT MEMORY ACCESS -----	205
8.5	I/O SYSTEM OPERATIONS -----	206
8.5.1	GENERAL -----	206
8.5.2	INITIALIZATION -----	206
8.5.3	COMMAND EXECUTION -----	208
8.5.4	DATA TRANSFERS -----	209
8.5.5	STATUS INTERRUPT HANDLING -----	211
8.5.6	DIRECT REQUEST COMMANDS -----	212
8.5.7	DATA CHAINING -----	213
8.6	EXAMPLE OF AN ACTUAL DEVICE CONTROLLER -----	214
8.6.1	DEVICE SPECIFICATIONS -----	214
8.6.2	CONTROLLER SPECIFICATIONS -----	215
8.6.2.1	DEVICE STATUS -----	215
8.6.2.2	DEVICE COMMANDS -----	217
8.6.2.3	DEVICE CONTROL WORDS -----	219
8.6.3	OPERATION AND PROGRAMMING -----	221
8.6.3.1	DATA TRANSFER OPERATION -----	222
8.6.3.2	DEVICE CONTROL FUNCTIONS -----	224
9	INSTALLATION PLANNING -----	225
9.1	GENERAL -----	225
9.2	SPACE REQUIREMENTS -----	225
9.3	ENVIRONMENTAL REQUIREMENTS -----	226
9.4	POWER REQUIREMENTS -----	226
9.5	PERIPHERAL REQUIREMENTS -----	227
9.6	TYPICAL INSTALLATION LAYOUT -----	228

APPENDICES

APPENDIX A - QM-1 PORT INTERFACE SPECIFICATIONS

APPENDIX B - QM-1 CPU OPTIONAL FEATURES

1. SCOPE AND PURPOSE

THE QM-1 IS A HIGH-SPEED GENERAL-PURPOSE DIGITAL COMPUTER THAT OPERATES UNDER TWO LEVELS OF MICROPROGRAM CONTROL. THE UNIQUE DESIGN OF THE QM-1 SUPPORTS A SYSTEM OF SOFTWARE-CREATED USER LEVELS, WHEREBY USERS AT DIFFERENT LEVELS APPROACH ARCHITECTURE, MACHINE LANGUAGE, AND PROGRAMMING IN WAYS MOST SUITED TO THEIR OWN SPECIFIC REQUIREMENTS OF THE HARDWARE. THE PRESENT DOCUMENT EXPLAINS THESE CONCEPTS AND DEFINES THE QM-1 AS IT APPEARS TO THE "HARDWARE-LEVEL" USER.

THE "HARDWARE-LEVEL" USER APPROACHES A PROGRAMMING INTERFACE WHOSE FUNCTIONAL PARTS CORRESPOND TO THE FACILITIES PROVIDED BY THE PHYSICAL QM-1 COMPUTER ITSELF, WITHOUT ANY RESTRICTIONS TO THE FULL GENERALITY OF THE HARDWARE IMPOSED BY PRE-DEFINITION OF THE CONTENTS OF ANY OF THE MACHINES CONTROL MEMORIES. EVEN THE CONTENTS OF THE READ-ONLY MEMORIES, INCLUDED FOR MACHINE BOOTSTRAP AND DIAGNOSTIC PURPOSES, MAY BE PROGRAMMED BY THIS USER. THE HARDWARE-LEVEL USER'S MANUAL, WHILE NOT AN ENGINEERING OR MAINTENANCE DOCUMENT, IS THUS ORIENTED TOWARD THE QM-1 USER WHOSE PURPOSE IS TO DEFINE HIS OWN COMPUTER STARTING AT THE LOWEST POSSIBLE LEVEL.

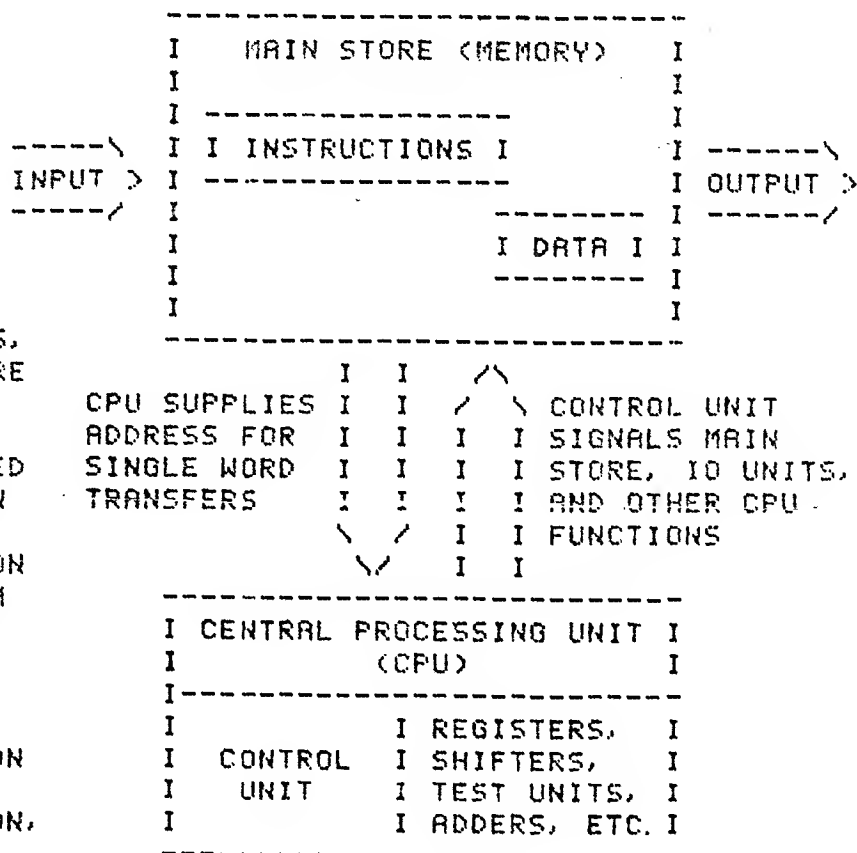
2. INTRODUCTION TO MICROPROGRAMMING AND THE QM-1

EVERY PROGRAMMABLE DEVICE, OR "MACHINE", POSSESSES AN ARCHITECTURE AND AN INSTRUCTION SET. THE ARCHITECTURE IS ITS SYSTEM OF COMPONENTS AND THEIR INTERCONNECTIONS; IN THE CASE OF A COMPUTER, ARCHITECTURES ARE DESCRIBED IN TERMS OF STORES, REGISTERS, ARITHMETIC-LOGIC UNITS, DATA PATHS, ETC. A MACHINE INSTRUCTION IS A COMMAND WHICH CAUSES ELEMENTS OF THE ARCHITECTURE TO OPERATE IN SOME PREDETERMINED MANNER; THE INSTRUCTION SET OF A MACHINE IS SIMPLY A LIST OF ALL INSTRUCTIONS WHICH THE MACHINE RECOGNIZES.

USING THESE BROAD DEFINITIONS AND THE SIMPLIFIED MODEL OF A COMPUTER SHOWN IN FIGURE 2A, A DISCUSSION OF THREE PHASES OF THE "INSTRUCTION SEQUENCE" PROVIDES A BASIC EXPLANATION OF COMPUTER OPERATION.

INSTRUCTION FETCH

SEQUENCES OF MACHINE INSTRUCTIONS, IN THE FORM OF BINARY NUMBERS, ARE TYPICALLY STORED IN CONTIGUOUS LOCATIONS IN MAIN STORE (MEMORY); INSTRUCTION EXECUTION IS INITIATED BY FETCHING A MACHINE INSTRUCTION FROM A GIVEN LOCATION IN MEMORY AND PLACING IT INTO AN INSTRUCTION REGISTER. THE MEMORY ADDRESS FROM WHICH TO FETCH AN INSTRUCTION IS CONTAINED IN AN INSTRUCTION LOCATION COUNTER REGISTER, OFTEN CALLED A PROGRAM COUNTER; PART OF THE EFFECT OF EVERY INSTRUCTION IS TO UPDATE THIS REGISTER TO POINT TO THE SUCCESSOR INSTRUCTION, AND THEN TO BEGIN THE MEMORY FETCH FOR THE NEXT SEQUENTIAL INSTRUCTION.



BLOCK DIAGRAM OF A COMPUTER
FIGURE 2A

INSTRUCTION DECODE

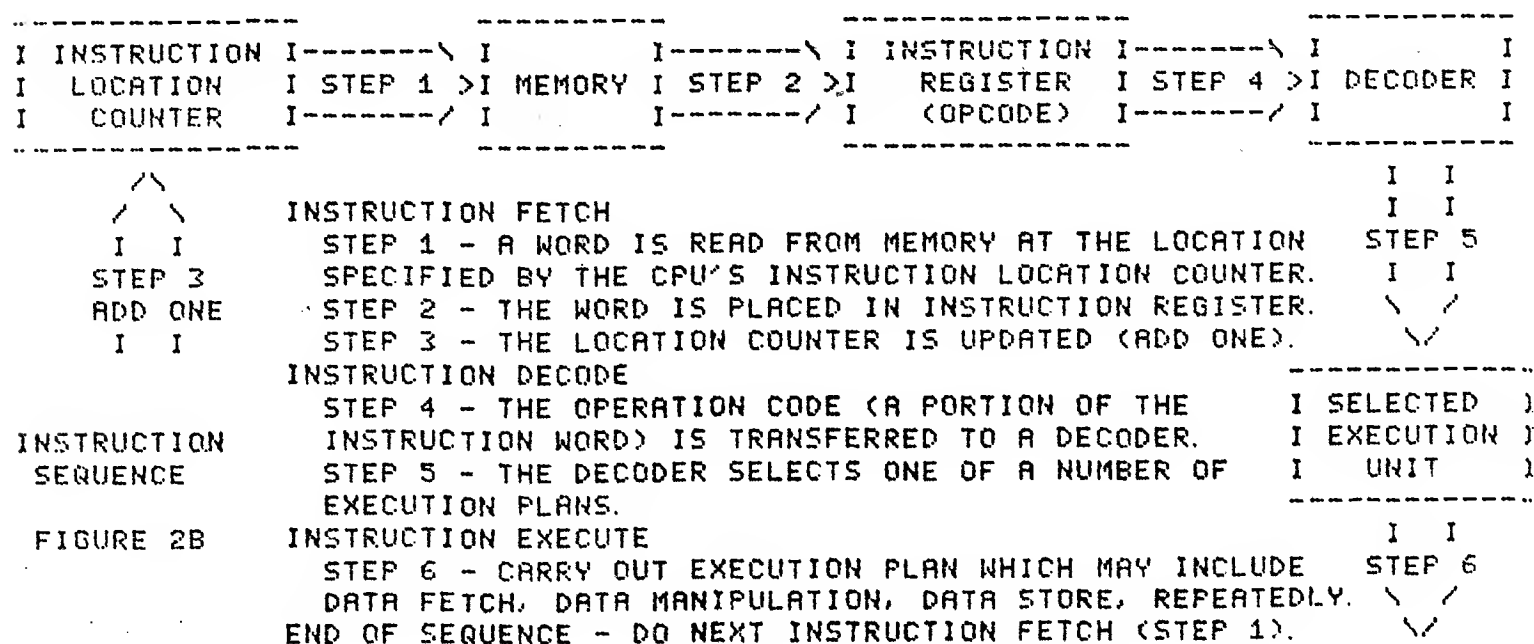
A PORTION OF THE CONTENTS OF THE INSTRUCTION REGISTER IS DESIGNATED AS THE OPERATION CODE. THIS BINARY NUMBER IS DECODED BY THE CONTROL UNIT TO SELECT AMONG A NUMBER OF MODULES, EACH OF WHICH IS RESPONSIBLE FOR ACCOMPLISHING THE EFFECT OF ONE OF THE INSTRUCTIONS IN THE COMPUTER'S INSTRUCTION SET. AS WILL BE SHOWN LATER, THE METHOD OF DECODING AND THE NATURE OF THESE MODULES IS CRITICAL TO THE DEFINITION OF MICROPROGRAMMING.

INSTRUCTION EXECUTION

THE ULTIMATE EFFECT OF ANY INSTRUCTION-EXECUTION MODULE IS THE GENERATION OF ELECTRICAL SIGNALS TO THE VARIOUS COMPUTER COMPONENTS.

2.1 BASIC INSTRUCTION SEQUENCE

THESE THREE PHASES OF INSTRUCTION FETCH, DECODE, AND EXECUTE, FORM THE BASIC INSTRUCTION SEQUENCE (OR "INSTRUCTION CYCLE"). AFTER INITIAL START-UP, ALL COMPUTERS FOLLOW AN INSTRUCTION SEQUENCE SIMILAR TO THAT ILLUSTRATED IN FIGURE 2B.



2.2 MICROPROGRAMMED CONTROL

THE FINAL PHASE IS OF PARTICULAR INTEREST HERE. THE ELECTRICAL SIGNALS WHICH THE CONTROL UNIT SENDS TO THE ARCHITECTURAL COMPONENTS ARE THE MOST BASIC, OR "PRIMITIVE", COMMANDS IN THE COMPUTER; THESE SIGNALS HAVE EFFECTS SUCH AS OPENING AND CLOSING GATES (FOR EXAMPLE, TO TRANSFER REGISTER CONTENTS), INITIATING MEMORY CYCLES, AND SETTING INDIVIDUAL BITS. IN FACT, THE INSTRUCTION SEQUENCE ITSELF IS UNDER THE CONTROL OF SUCH PRIMITIVE OPERATIONS; AN IMPLICIT EFFECT OF EVERY MACHINE INSTRUCTION IS THE EXECUTION OF THE NEXT INSTRUCTION SEQUENCE.

ONLY RARELY DO MACHINE INSTRUCTIONS CORRESPOND TO A SINGLE ARCHITECTURAL PRIMITIVE; MOST MACHINE INSTRUCTIONS RESULT IN THE GENERATION OF A NUMBER OF PRIMITIVES, FREQUENTLY ARRANGED IN A TIME SEQUENCE. FOR SOME INSTRUCTIONS, THE ARRANGEMENT OF PRIMITIVES CAN BE FAIRLY COMPLEX. AN EXAMPLE IS A MULTIPLY INSTRUCTION ON A MACHINE WHICH HAS ONLY A ADDING COMPONENT; THE ADDER MUST BE USED ITERATIVELY, AND THE INTERNAL PLAN OF THE INSTRUCTION RESEMBLES A COMPUTER PROGRAM.

THE LATER OBSERVATION SUGGESTS AN IMPLEMENTATION OF THE PRIMITIVE SIGNAL CONTROL FUNCTION. IN THE CONVENTIONAL, OR "HARD-WIRED" COMPUTER, A HARDWARE DECODING OF THE RELEVANT PORTION OF THE INSTRUCTION WORD SELECTS ONE OF SEVERAL LOGIC CIRCUITS, EACH OF WHICH IS RESPONSIBLE FOR GENERATING AND SEQUENCING THE PRIMITIVE SIGNALS OF A GIVEN MACHINE INSTRUCTION. IF, HOWEVER, THE PRIMITIVE CONTROL FUNCTIONS ARE REGARDED AS "MICRO-OPERATIONS", THEN A "MICROPROGRAM" CAN BE WRITTEN TO PLAN THE FLOW OF AN INSTRUCTION. THE STEPS OF THIS MICROPROGRAM CAN THEN BE IMPLEMENTED AS PRIMITIVE COMMANDS EXECUTING OUT OF A FAST-ACCESS STORE, SUCH AS SEMICONDUCTOR MEMORY. (EXECUTION OF SUCH COMMANDS IS SIMPLE TO ACCOMPLISH, SINCE THE MICRO-OPERATIONS CORRESPOND DIRECTLY TO ARCHITECTURAL FUNCTIONS.)

USING A MICROPROGRAMMED APPROACH TO MACHINE INSTRUCTION IMPLEMENTATION, THE INSTRUCTION-DECODE STEP OF MACHINE OPERATION CHANGES: RATHER THAN DECODING THE OPERATION-CODE PORTION OF THE INSTRUCTION TO SELECT ONE OF SEVERAL HARDWARE MODULES, THIS BINARY NUMBER IS USED DIRECTLY AS AN ADDRESS, OR POINTER, INTO THE MICROPROGRAM STORE ("CONTROL STORE"); THE LOCATION SO DEFINED IS PROGRAMMED AS THE ENTRY POINT OF THE MICROPROGRAM WHICH IMPLEMENTS THE ORIGINAL MACHINE INSTRUCTION. THIS PROCESS IS ILLUSTRATED IN FIGURE 2C.

2.3 USES OF MICROPROGRAMMING

WITH THE PREVIOUSLY DEFINED MODEL OF MICROPROGRAM MACHINE CONTROL, WE CAN NOW EXAMINE THE USES AND ADVANTAGES OF MICROPROGRAMMING. THE STRONGEST SINGLE JUSTIFICATION FOR MICROPROGRAMMING LIES IN THE CURRENT DISPARITY BETWEEN THE SPEED OF MAIN STORE (CORE MEMORY) AND THE SPEED OF CURRENTLY AVAILABLE LOGICAL COMPONENTS. FOR EXAMPLE, MORE THAN 10 SETS OF PRIMITIVE FUNCTIONS MAY BE EXECUTED IN THE TIME TAKEN TO READ ONE WORD FROM CORE MEMORY. THUS TIME EXISTS FOR MORE THAN 10 CONTROL STORE STEPS TO IMPLEMENT A MAIN STORE INSTRUCTION. THIS LARGE RATIO MAKES POSSIBLE A SIGNIFICANT INCREASE IN THE POWER OF THE INSTRUCTIONS DEFINED AT THE HIGHER LEVEL OVER THOSE REQUIRED IN THE UNDERLYING HARDWARE. FOR THIS REASON, MICROPROGRAMMING IS NOW COMMON IN MANY COMPUTERS.

MICROPROGRAMMING PROVIDES OTHER ADVANTAGES AS WELL. SINCE MICROPROGRAMMING IN CONTROL STORE SERVES TO DEFINE THE COMPUTER AS SEEN AT THE CONVENTIONAL LEVEL, THE FLEXIBILITY OF MICROPROGRAMMING MAY BE USED TO VARY THE MACHINE DEFINED. MANY OF THE ADVANTAGES THAT RESULT ARE TABULATED IN FIGURE 2D.

MICROPROGRAMMING MAY BE USED TO	ADVANTAGES	FIGURE 2D
1. DEFINE A COMPUTERS INSTRUCTION SET INDEPENDENT OF THE BASIC HARDWARE DEVELOPMENT. THIS WAS THE MOST COMMON EARLY USE.	1 A) SEPERATES THE INSTRUCTION DEFINITION FROM THE HARDWARE SPECIFICATION. 1 B) PERMITS MATCHING MEMORY SPEED TO LOGIC SPEEDS WHEN A LARGE DIFFERENCE EXISTS.	
2. CAUSE THE HARDWARE TO FUNCTION AS ANOTHER (PRE-EXISTING) COMPUTER. THIS IS THE COMMON DEFINITION OF EMULATION.	1 A) EMULATED MACHINES SOFTWARE MAY BE USED WITHOUT MODIFICATION THUS PRESERVING POSSIBLY LARGE SOFTWARE INVESTMENTS. 1 B) HOST COMPUTER SYSTEM MAY BE FASTER OR LESS EXPENSIVE THAN ORIGINAL MACHINE. 1 C) SEVERAL MACHINES MAY BE EMULATED AT DIFFERENT TIMES, ON SAME HARDWARE HOST.	
3. EMULATE ANOTHER COMPUTER, BUT WITH EXTRA INSTRUCTIONS AND/OR SPECIAL FEATURES.	1 INCREASED EFFICIENCY: FUNCTIONS REQUIRING COMPLEX AND TIME CONSUMING SOFTWARE MAY BE PERFORMED DIRECTLY ON THE MACHINE, AS A SINGLE (SPECIAL) INSTRUCTION. EXAMPLES: 1 * FLOATING POINT ARITHMETIC 1 * OPERATING SYSTEM FUNCTIONS 1 * ANY PROGRAMMED PROCEDURE COMMONLY USED IN A GIVEN APPLICATION.	

-
- | | | |
|---|---|---|
| 4. CREATE A SPECIAL-PURPOSE
COMPUTER TO MEET THE NEEDS OF
A PARTICULAR ENVIRONMENT. | I | A) MICROPROGRAM DEVELOPMENT IS EASIER,
FASTER, AND LESS EXPENSIVE THAN
HARDWARE DEVELOPMENT, AND IS PERFORMED
BY PERSONNEL TYPICALLY CLOSER TO END
NEEDS THAN HARDWARE PERSONNEL. |
| | I | B) RESULT CAN BE MODIFIED EASILY WHEN
NECESSARY, AS NEEDS CHANGE. |
| | I | C) WHEN APPLICATION IS PHASED OUT, HOST
HARDWARE REMAINS USABLE. |
-

SINCE A FULLY FLEXIBLE MICROPROGRAMMED DESIGN ALSO PERFORMS THE INSTRUCTION FETCH AND INSTRUCTION DECODE UNDER MICROPROGRAM CONTROL, MAIN STORE BECOMES MERELY A STORAGE AREA WHICH MAY (AMONG OTHER THINGS) CONTAIN INSTRUCTIONS OF A HIGHER-LEVEL MACHINE. THEREFORE A FIFTH USE OF MICROPROGRAMMING IS:

- | | | |
|--|---|---|
| 5. WRITE USER PROGRAMS IN CONTROL
STORE, WITH MAIN STORE USED AS
A FAST MESSAGE BUFFER, PAGE
BACKUP, FILE STORAGE, ETC. | I | VERY FAST PROCESSING TIMES ARE POSSIBLE
FOR SUITABLE APPLICATIONS. LESS
HARDWARE MAY BE NECESSARY TO DO THE JOB
SINCE THE HARDWARE IS USED DIRECTLY. |
|--|---|---|
-

2.4 HORIZONTAL AND VERTICAL CONTROL

THE DESIGNER OF A MACHINE WITH MICROPROGRAMMED CONTROL FACES AN IMMEDIATE DECISION AS TO THE FORMAT OF MICROINSTRUCTIONS TO BE USED IN THE MACHINE. HE MAY CHOOSE TO USE A WIDE, UNSTRUCTURED MICROWORD, USUALLY CALLED A HORIZONTAL MICROINSTRUCTION:

```

-----
[.....HORIZONTAL MICROINSTRUCTION.....]
-----

```

EACH BIT IS
INDEPENDENT
OF OTHER BITS.

WHEN EXECUTED, EACH BIT IN A HORIZONTAL MICROINSTRUCTION RESULTS IN A CONTROL SIGNAL TO A HARDWARE COMPONENT. THIS IS GENERALLY FOUND IN MORE POWERFUL MACHINES. THE MICROINSTRUCTION MAY RUN TO 100 OR MORE BITS (THE IBM 360/50 USES A MICROINSTRUCTION 90 BITS WIDE).

OR THE DESIGNER MAY CHOOSE A HIGHLY ENCODED MICROINSTRUCTION PACKED INTO A MUCH SMALLER WORD. THE WORD CONTAINS A MICRO-OPCODE AND SEVERAL OTHER ENCODED FIELDS. FOR THIS REASON, IT IS OFTEN REFERRED TO AS A VERTICAL MICROINSTRUCTION:

```

VERTICAL MICROINSTRUCTION      [ MICRO-OPCODE /  XXX XXXX ]
-----

```

TOGETHER, SEVERAL
BITS FORM AN
ENCODED FIELD.

WHEN EXECUTED, THE MICRO-OPCODE OF A VERTICAL MICROINSTRUCTION SELECTS A SEQUENCE OF CONTROL SIGNALS, SIMILAR TO THE OPERATION OF A MACHINE INSTRUCTION OPCODE BUT AT A LOWER LEVEL (SIMPLER SEQUENCES ARE INVOKED). VERTICAL MICROINSTRUCTIONS ARE MUCH SHORTER (THE IBM 360/25 HAS A 16 BIT MICROINSTRUCTION).

EACH SCHEME FOR MICROPROGRAMMED CONTROL OFFER CERTAIN ADVANTAGES. A CHOICE INVOLVES EVALUATION OF MANY TRADE-OFFS. SOME OF THE FACTORS ARE TABULATED IN FIGURE 2E.

TRADE-OFFS BETWEEN HORIZONTAL AND VERTICAL CONTROL

FIGURE 2E

HORIZONTAL MICROINSTRUCTIONS...

VERTICAL MICROINSTRUCTIONS...

-----I-----	I	-----I-----
ALLOW ULTIMATE FLEXIBILITY IN	I	PROVIDE A LIMITED SELECTION OF
CONTROL, SINCE EACH SIGNAL (BIT)	I	CONTROL PATTERNS; THE NUMBER OF

MAY BE INDIVIDUALLY SELECTED BY THE MICROPROGRAMMER.	I I	POSSIBILITIES DEPENDS UPON THE WIDTH OF THE MICRO-OPCODE.
-----	I	-----
MAY BE EXECUTED SIMPLY BY GATING THEM TO A REGISTER, TO WHICH SIGNAL LINES ARE ATTACHED DIRECTLY.	I I I I	REQUIRE EXECUTION MACHINERY SIMILAR TO (BUT SIMPLER THAN) THAT REQUIRED TO EXECUTE MACHINE INSTRUCTIONS.
-----	I	-----
ALLOW PARALLEL OPERATION OF HARDWARE COMPONENTS.	I I	TYPICALLY SPECIFY "SINGLE-THREAD" OPERATIONS.
-----	I	-----
ARE RELATIVELY DIFFICULT TO PROGRAM.	I I	ARE RELATIVELY SIMPLE TO PROGRAM.
-----	I	-----
MUST BE EXECUTED FREQUENTLY, SINCE THEY EXERCISE EACH HARDWARE COMPONENT AT MOST ONCE.	I I I	MAY SPECIFY A TIME-SEQUENCE OF CONTROL SIGNALS, SO THEY MAY BE EXECUTED LESS FREQUENTLY.
-----	I	-----
ARE NIDE, TYPICALLY ON THE ORDER OF 100 BITS.	I I I I	REQUIRE ONLY ENOUGH BITS TO CONTAIN THE MICRO-OPCODE AND PERHAPS SOME PARAMETERS -- TYPICALLY 8 TO 16 BITS.
-----	I	-----
THE LAST TWO ITEMS IMPLY THAT STORAGE OF ENOUGH HORIZONTAL MICROINSTRUCTIONS TO RUN A REASONABLY POWERFUL EMULATION MAY BE PROHIBITIVELY EXPENSIVE IN NUMBER OF BITS.	I I I I I I I	THE LAST TWO ITEMS IMPLY THAT STORAGE OF ENOUGH VERTICAL MICROINSTRUCTIONS TO RUN A REASONABLY POWERFUL EMULATION MAY BE ACCEPTABLY INEXPENSIVE IN NUMBER OF BITS.
-----	I	-----

CONCLUSIONS:

HORIZONTAL MICROINSTRUCTIONS ARE PREFERABLE TO VERTICAL MICROINSTRUCTIONS FOR FLEXIBILITY AND PARALLELISM, BUT THEY ARE MORE DIFFICULT TO PROGRAM, REQUIRE LARGER AMOUNTS OF EXPENSIVE STORAGE, AND ARE LIMITED IN WHAT TIME SEQUENCES MAY BE PROGRAMMED.

THE QM-1 HAS BEEN DESIGNED TO MAKE AVAILABLE THE ADVANTAGES OF EACH SCHEME OF MICROPROGRAMMED CONTROL AND TO AVOID THE DISADVANTAGES INHERENT IN EACH. THE UNIQUE FEATURES OF THE QM-1 THAT MAKE THIS POSSIBLE WILL BE EXAMINED NEXT.

2.5 THE QM-1 CONTROL HIERARCHY

IN THE QM-1, A TWO-LEVEL DESIGN SMOOTHS THE MACHINE DEFINITION PROCESS OVER TWO STAGES, ACHIEVING THE ADVANTAGES OF BOTH HORIZONTAL AND VERTICAL CONTROL:

MACHINE INSTRUCTIONS IN MAIN STORE ARE EXECUTED BY (AND DEFINED BY) MICROPROGRAMS IN CONTROL STORE, UNDER VERTICAL CONTROL.

MICROINSTRUCTIONS IN CONTROL STORE ARE IN TURN EXECUTED BY (AND DEFINED BY) NANOPROGRAMS IN NANOSTORE, UNDER HORIZONTAL CONTROL.

AN ILLUSTRATION OF THIS CONCEPT IS SHOWN IN FIGURE 2F - QM-1 CONTROL HIERARCHY.

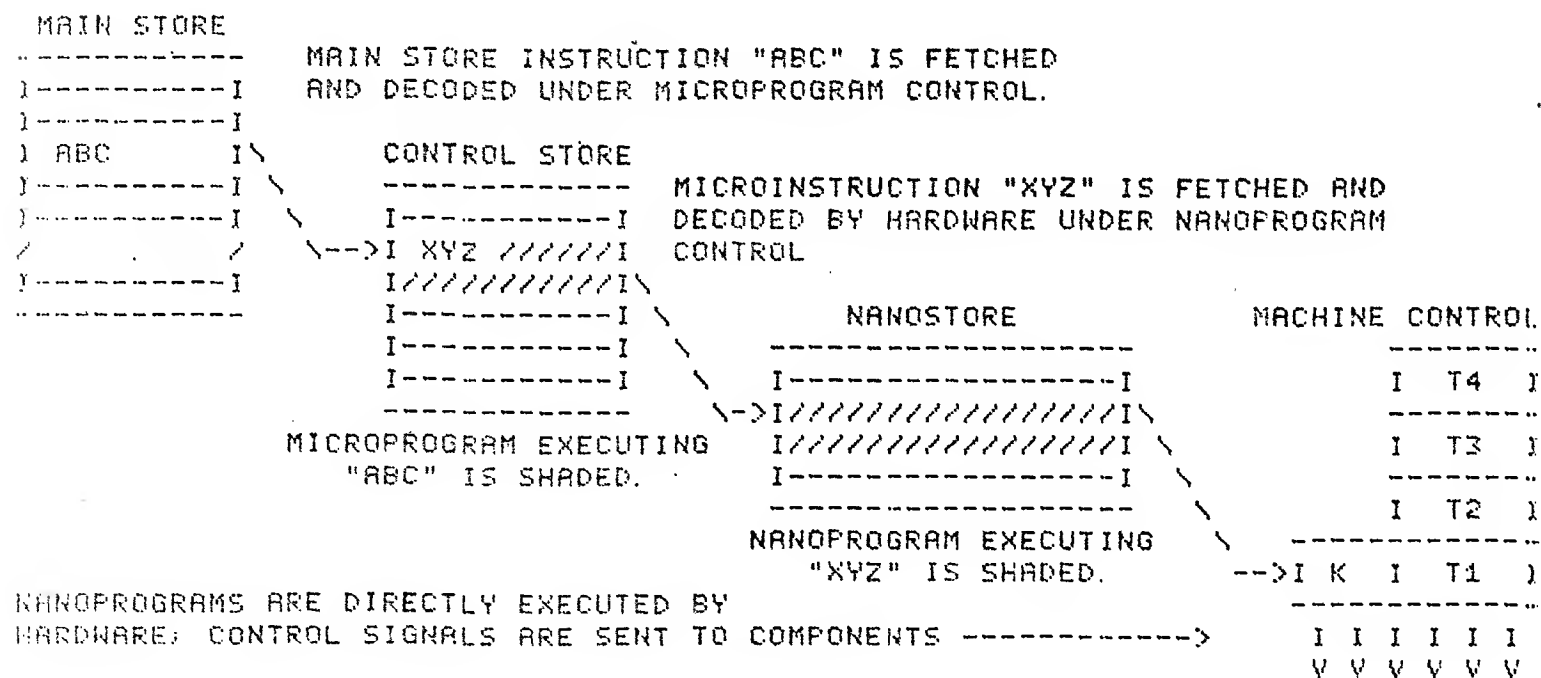


FIGURE 2F

QM-1 CONTROL HIERARCHY

EXAMPLE OF TWO LEVEL EMULATION

THIS UNIQUE CONTROL HIERARCHY TAKES ADVANTAGES OF THE BEST FEATURES OF BOTH HORIZONTAL AND VERTICAL CONTROL AS SUMMARIZED IN FIGURE 2G. IN ADDITION, FLEXIBLE TIME SEQUENCING IS POSSIBLE AT BOTH LEVELS. AND MOST IMPORTANT, BOTH CONTROL STORE AND NANOSTORE ARE FULLY WRITABLE SEMICONDUCTOR MEMORIES, SO THAT THE QM-1 USER CAN TAKE ADVANTAGE OF ALL POSSIBLE FLEXIBILITY IN THE SYSTEM BY DYNAMIC REPROGRAMMING.

IN PARTICULAR, CONTROL STORE IS A FULLY GENERAL-PURPOSE READ/WRITE STORE; HENCE IT IS FEASIBLE, FOR SOME APPLICATIONS, TO APPROACH QM-1 CONTROL STORE AS THE PRIMARY PROGRAM STORE OF THE MACHINE, EXECUTING PROGRAMS WHICH CAN REGARD THE PASSIVE MAIN STORE AS A SECONDARY STORAGE UNIT.

CONTROL HIERARCHY DIMENSIONAL TRADE-OFFS

FIGURE 2G

AT HIGHEST LEVEL

END USER HAS SYSTEM
SIMPLE TO PROGRAM. <----->

GENERALIZED INDIRECT CONTROL
POWERFUL (HIGH LEVEL) <----->
INSTRUCTIONS

MEANING OF MAIN STORE
CONTENTS FULLY REDEFINABLE <----->

LARGE MEMORY AVAILABLE <----->

LOW COST/BIT <----->

AT LOWEST LEVEL

HARDWARE DESIGNER HAS SYSTEM
DIRECT TO IMPLEMENT.

ABSOLUTE DIRECT CONTROL.
PRIMITIVE (LOW LEVEL)
FUNCTIONS.

MEANING OF CONTROL
SIGNALS FIXED IN HARDWARE.

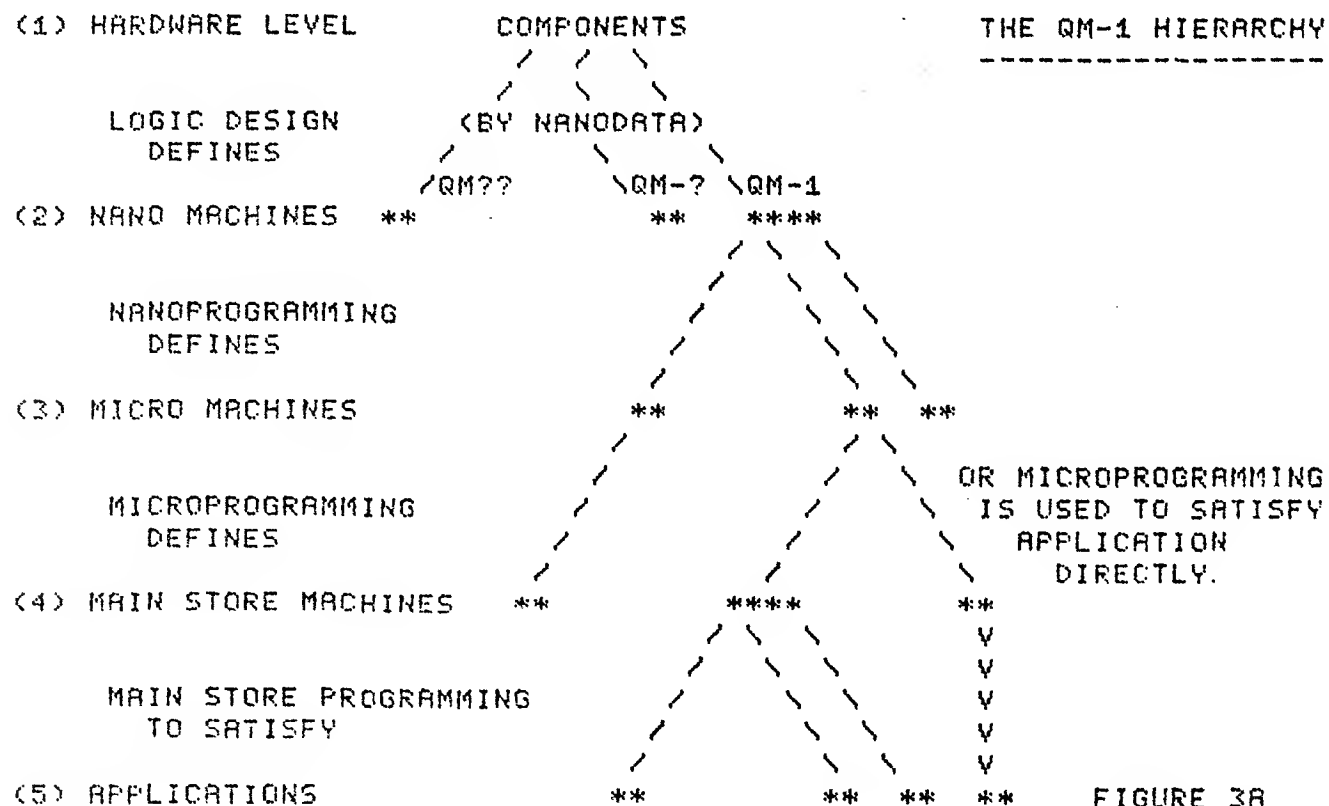
SMALL STORE REQUIRED.

FAST OPERATION.

3. USER AND MACHINE HIERARCHIES

THE DESIGN OF THE QM-1 SUGGESTS THE USE OF A SYSTEM OF "VIRTUAL MACHINES" ARRANGED IN A HIERARCHY OF LEVELS. EACH LEVEL IS SUPPORTED BY THE MACHINE BELOW, AND IN TURN SUPPORTS THE MACHINE ABOVE. ONCE A GIVEN MACHINE IS DEFINED BY SUITABLE SOFTWARE (OR "FIRMWARE"), ITS IMPLEMENTATION -- I.E., THE NATURE OF THAT SOFTWARE STRUCTURE -- IS TRANSPARENT TO THE USER OF THAT MACHINE. FOR EXAMPLE, AFTER SUITABLE NANOPROGRAMMING IS DONE TO DEFINE A "MICRO-MACHINE", THE VERY EXISTENCE OF NANOSTORE IS IRRELEVANT TO THE MICRO-MACHINE USER.

SUCH A MACHINE HIERARCHY IS SHOWN IN FIGURE 3A AND DESCRIBED IN DETAIL BELOW.



HARDWARE LEVEL

THE BASIC HARDWARE COMPONENTS OF THE QM-1 INCLUDE SEVERAL BANKS OF REGISTERS; A SYSTEM OF THREE STORES; ARITHMETIC, BOOLEAN, AND SHIFT COMPONENTS; AND TWELVE INDEPENDENT BUSES. BUS CONNECTIONS BETWEEN THE COMPONENTS ARE PROGRAMMABLE AND MAY BE CHANGED AS OFTEN AS REQUIRED TO BEST FIT THE CURRENT TASK. ALL THESE UNITS MAY BE EXERCISED INDEPENDENTLY, ALLOWING A HIGH DEGREE OF PARALLELISM.

COMPLETE CONTROL OVER THE HARDWARE IS PROVIDED BY A 360-BIT WORD READ FROM THE DYNAMICALLY WRITABLE NANOSTORE; THE ACTIVE NANOWORD PROVIDES A SEQUENCE OF FOUR MACHINE STATE VECTORS, EACH OF WHICH DRIVES THE INDIVIDUAL MACHINE COMPONENTS AND THEIR INTERCONNECTIONS DURING A MACHINE CLOCK PERIOD OF 60 NANOSECONDS.

NANO-MACHINE LEVEL

NANOPROGRAMMING IS THE PROCESS OF DEFINING A SET OF SUCH CONTROL SEQUENCES TO IMPLEMENT MICROINSTRUCTIONS EXECUTED AT THE NEXT LEVEL. THE OPCODE OF A VERTICALLY FORMATTED MICROINSTRUCTION, READ FROM CONTROL STORE, IS USED TO SELECT THE ENTRY POINT IN NANOSTORE AT WHICH TO BEGIN EXECUTING THE DEFINING NANOPROGRAM. THE MICROINSTRUCTION SET USED MAY BE EITHER THAT DEFINED BY NANODATA (WITH POSSIBLE USER MODIFICATIONS/EXTENSIONS FOR THE CURRENT TASK) OR THAT DEFINED BY THE USER; THE NANODATA-SUPPLIED MICROLANGUAGE IS ACCOMPANIED BY SYSTEMS SOFTWARE TO SUPPORT IO AND PROCESS MANAGEMENT.

MICRO-MACHINE LEVEL

SINCE MICROINSTRUCTIONS RESIDE IN THE FULLY READABLE/WRITABLE CONTROL STORE, MICROPROGRAMMING CAN BE USED TO DEFINE THE APPLICATION DIRECTLY. DUE TO THE FLEXIBILITY PROVIDED AT THE NANO LEVEL, A VARIETY OF MICRO-MACHINES MAY BE DEFINED TO EFFICIENTLY MATCH THE APPLICATION. THE MICRO-MACHINE CAN THEN BE VIEWED AS A CONVENTIONAL MACHINE WITH A CUSTOMIZED INSTRUCTION SET AND A 120-NANOSECOND MEMORY.

MAIN-STORE-MACHINE LEVEL

FOR MANY APPLICATIONS, THE ABOVE NUMBER OF LEVELS WILL BE SUFFICIENT; APPLICATIONS SOFTWARE MAY BE WRITTEN IN THE DEFINED MICROLANGUAGE, EXECUTING OUT OF CONTROL STORE AT VERY HIGH SPEEDS. FOR THOSE APPLICATIONS IN WHICH ANOTHER LEVEL OF FLEXIBILITY IS DESIRED, HOWEVER, MICROPROGRAMMING IN CONTROL STORE MAY BE USED TO DEFINE THE ARCHITECTURE AND INSTRUCTION SET FOR SOFTWARE IN MAIN STORE. AT THE MICRO LEVEL, MAIN STORE IS VIEWED SIMPLY AS A PASSIVE GENERAL-PURPOSE DATA STORE; THE PROCESS IS ONE OF CLASSICAL EMULATION.

AS INDICATED IN SECTION 1, THE PURPOSE OF THIS MANUAL IS TO PROVIDE COMPLETE FUNCTIONAL SPECIFICATIONS OF THE QM-1, AND THUS TO DEFINE THE "NANO-MACHINE" AVAILABLE TO THE HARDWARE-LEVEL USER. MANY USERS WILL BE CONCERNED WITH THE MACHINE AT THIS MOST FUNDAMENTAL LEVEL. THE NANODATA SYSTEMS SOFTWARE STAFF, FOR EXAMPLE, APPROACHES THE MACHINE AT THIS LEVEL.

WHEN APPROPRIATE SOFTWARE, INCLUDING BOTH SYSTEMS SUPPORT FUNCTIONS AND ANY ONE OF SEVERAL MICRO LANGUAGE DEFINITIONS, IS INCLUDED IN THE QM-1, THE MICRO-LEVEL USER CAN PROGRAM THE MACHINE WITHOUT BEING CONCERNED WITH THE STRUCTURE BENEATH.

THUS THIS MANUAL IS DEDICATED TO THAT NEW BREED - THE NANOPROGRAMMER. OTHER PROGRAMMERS MAY HAVE INTEREST IN THE MANUAL IN ORDER TO UNDERSTAND THE HARDWARE THAT SUPPORTS THE LEVEL AT WHICH THEY WRITE PROGRAMS; THE HARDWARE-LEVEL USER WILL FIND THE MATERIAL IN THE NEXT TWO SECTIONS ESSENTIAL.

4. QM-1 FUNCTIONAL SPECIFICATIONS, PART I

4.1 GENERAL

SECTIONS 4 AND 5 OF THIS MANUAL ARE A COMPLETE FUNCTIONAL SPECIFICATION OF THE QM-1 CPU IN TWO PARTS. PART II (SECTION 5) IS INTENDED TO BE USED AS A PROGRAMMER'S REFERENCE GUIDE, AND INCLUDES CONTROL FIELD MAPPINGS AND ENCODINGS AS WELL AS DETAILED FUNCTIONAL DESCRIPTION. PART I (SECTION 4) EXPLAINS QM-1 MACHINE CONCEPTS, ARCHITECTURE, AND OPERATIONS, AND PROVIDES AN OVERVIEW OF THE QM-1 AND AN INTRODUCTION TO PART II.

IN ORDER TO INTRODUCE THE MACHINE SPECIFICATIONS TO THE FIRST-TIME READER, PART I BECOMES PROGRESSIVELY MORE SPECIFIC AS MORE OF THE OVERALL QM-1 STRUCTURE IS REVEALED. HENCE THE EARLIER SECTIONS OF PART I HAVE MORE DETAILED EXPLANATIONS IN PART II.

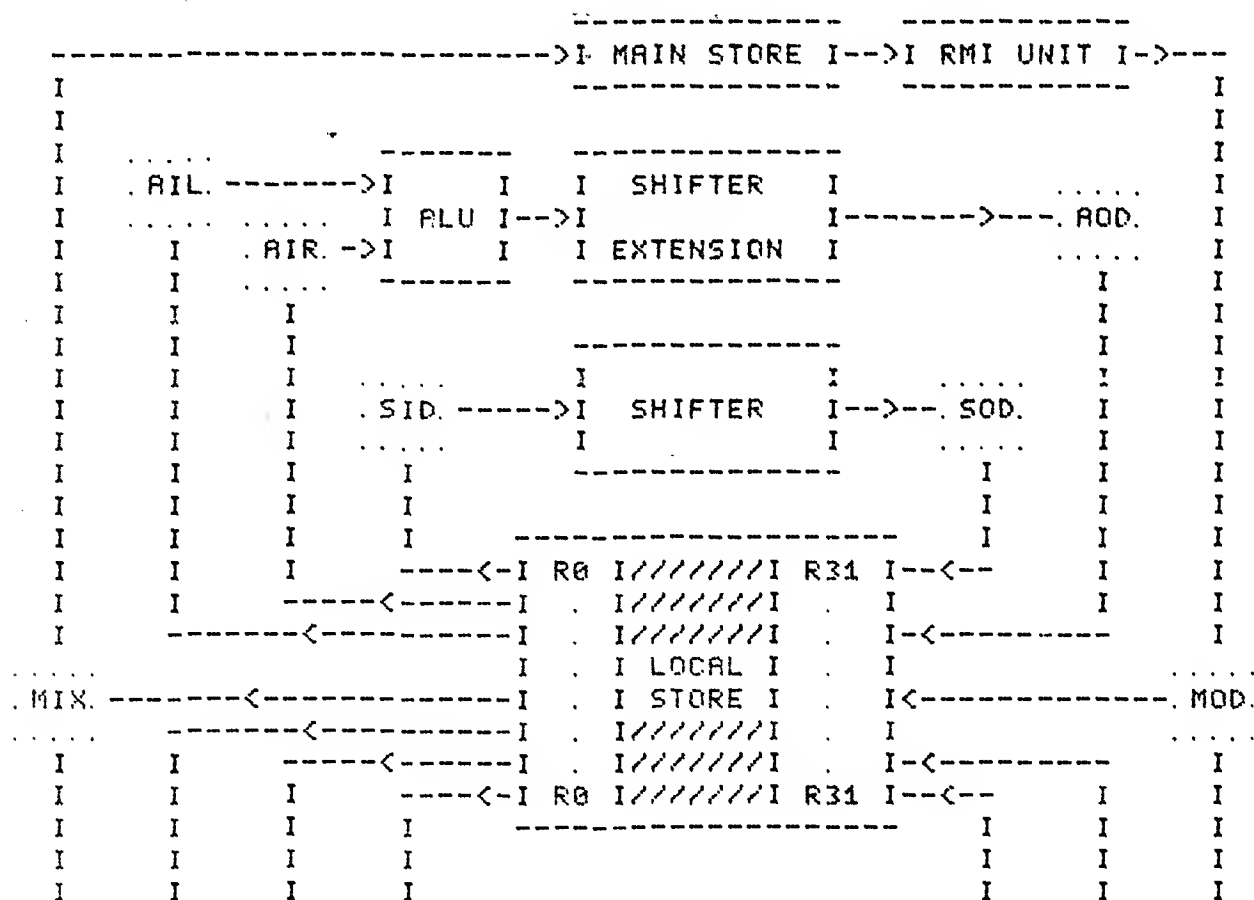
THE MACHINE DESCRIBED IN SECTIONS 4 AND 5 (QM-1 FUNCTIONAL SPECIFICATIONS) IS THE "HARDWARE QM-1". ITS ARCHITECTURAL FEATURES AND CONTROLS ARE THOSE AVAILABLE TO THE LOWEST LEVEL (NANO-) PROGRAMMER. NANOPROGRAMMING MAY BE USEFULLY VIEWED AS THE TASK OF IMPLEMENTING A ("VIRTUAL") MACHINE DEFINITION FOR USE AT THE NEXT HIGHER (MICRO) LEVEL. THE DESCRIPTION OF THE "MICRO-MACHINE" WILL NOT NECESSARILY RESEMBLE THAT OF THE QM-1, AND IN FACT MAY BE QUITE DIFFERENT. ANY NUMBER OF THE QM-1 RESOURCES MAY BE DEDICATED TO THE IMPLEMENTATION OF THE MICRO-MACHINE DEFINITION. THE NANOPROGRAMMER CAN ASSIGN SEVERAL OF THE 32 GENERAL-PURPOSE REGISTERS (LOCAL STORE) AVAILABLE TO HIM AS SPECIAL-PURPOSE ARCHITECTURAL FEATURES OF THE MICRO-MACHINE, E. G., INSTRUCTION REGISTER AND LOCATION COUNTER. THE MICRO-MACHINE ARCHITECTURE WILL IN GENERAL BE AN EXTENSION (RATHER THAN A RESTRICTION) OF THE QM-1 ARCHITECTURE; FOR EXAMPLE, A MICRO-MACHINE MAY BE DESIGNED WITH A LARGE NUMBER OF GENERAL-PURPOSE REGISTERS (WHICH THE NANOPROGRAMMER WOULD PROBABLY MAP INTO CONTROL STORE). THE RANGE OF FEASIBLE MICRO-MACHINE DEFINITIONS IS LIMITED ONLY BY THE INGENUITY OF THE NANOPROGRAMMER AND THE EFFICIENCY CONSIDERATIONS OF THE EMULATION PROCESS; STACK-MACHINE ARCHITECTURES, SOPHISTICATED ARITHMETIC PROCESSORS, AND "WIDE-WORD" MACHINES ARE, FOR EXAMPLE, WELL WITHIN THIS RANGE.

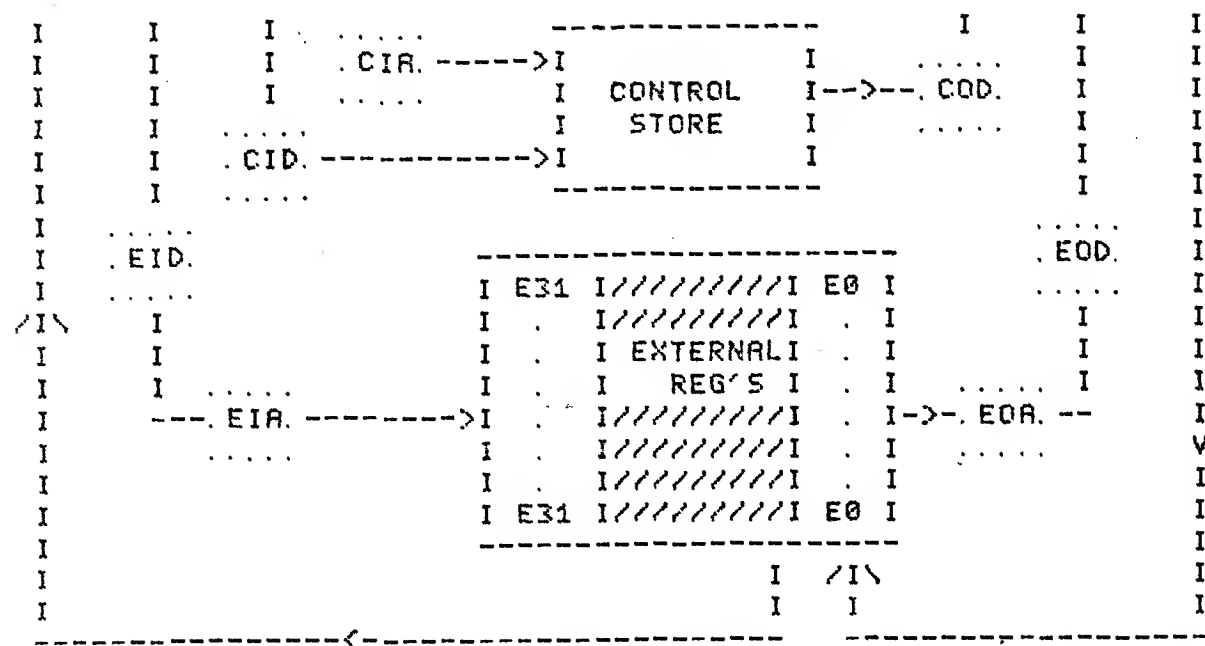
4.2 MAJOR RESOURCES AND ORGANIZATION

4.2.1 MAJOR BUSING STRUCTURE AND LOCAL STORE

THE MAJOR HARDWARE UNITS OF THE QM-1 -- STORES, ARITHMETIC-LOGIC UNIT, SHIFTER, REGISTER BANKS -- CAN EACH PROCESS OR STORE 18 BITS OF DATA IN PARALLEL, AND ARE CONNECTED BY A SYSTEM OF TWELVE 18-BIT-WIDE DATA PATHS (BUSES). THE CENTRAL MAJOR UNIT, LOCAL STORE, IS A TERMINUS FOR ALL TWELVE BUSES; THE OTHER END OF EACH BUS IS CONNECTED TO SOME OTHER MAJOR UNIT. THIS STRUCTURE IS SHOWN IN FIGURE 4.2.18.

QM-1 MAJOR BUSING STRUCTURE (18 BIT WIDE ARCHITECTURE) FIGURE 4.2.1A





ASSOCIATED WITH EACH BUS IS A DIRECTION OF DATA FLOW AND, IN GENERAL, A DISTINCT NANOPRIMITIVE CONTROL FOR THE GATING (TRANSMISSION) OF DATA. SINCE THE BUSES AND THEIR CONTROLS ARE PHYSICALLY INDEPENDENT, THEY MAY BE EXERCISED IN PARALLEL, ALLOWING A MAXIMUM OF TWELVE PROGRAM-CONTROLLED 18-BIT BUS TRANSFERS TO OCCUR SIMULTANEOUSLY. EACH BUS BEARS A THREE-LETTER LABEL STRUCTURED AS FOLLOWS:

THE FIRST LETTER CODES THE
MAJOR UNIT WHICH THE BUS
CONNECTS TO LOCAL STORE:

- M - MAIN STORE
- C - CONTROL STORE
- A - ARITHMETIC-LOGIC UNIT
(AND HIGH-ORDER HALF
OF SHIFTER OUTPUT)
- S - SHIFTER (LOW ORDER
18 BITS ONLY)
- E - EXTERNAL STORE

THE SECOND LETTER
DEFINES THE DIRECTION
OF DATA FLOW:

```

I - INPUT (TO THE
    NAMED UNIT FROM
    LOCAL STORE)

O - OUTPUT (FROM THE
    NAMED UNIT TO
    LOCAL STORE)

```

THE THIRD LETTER PROVIDES
FURTHER DESCRIPTIVE
INFORMATION:

```
A - ADDRESS
D - DATA
L - LEFT
R - RIGHT
X - "MULTIPLEX" (USED
    ONLY FOR MIX WHICH IS
    SHARED FOR MAIN STORE
    ADDRESS AND DATA)
```

THUS THE TWELVE BUSES ARE LABELED:

MIX - MAIN STORE INPUT - ADDRESS/DATA	MOD - MAIN STORE OUTPUT DATA
CIA - CONTROL STORE INPUT - ADDRESS	
CID - CONTROL STORE INPUT - DATA	COD - CONTROL STORE OUTPUT DATA
AIL - ALU INPUT - LEFT	
AIR - ALU INPUT - RIGHT	AOD - ALU OUTPUT DATA
SID - SHIFTER INPUT DATA	SOD - SHIFTER OUTPUT DATA
EID - EXTERNAL STORE INPUT DATA	EOD - EXTERNAL STORE OUTPUT DATA

EXPLANATION OF THE EOA AND EIA LABELS IN THE DIAGRAM IS DEFERRED TO SECTION 4.2.5.

MUCH OF THE PROGRAMMER'S ATTENTION INVOLVES THE INTERACTION OF LOCAL STORE WITH THE OTHER MAJOR UNITS, VIA THE BUSING STRUCTURE. LOCAL STORE IS A BANK OF 32 18-BIT REGISTERS, LOGICALLY UNIFORM WITH RESPECT TO BUSING. EACH BUS IS INDEPENDENTLY CONNECTABLE, UNDER PROGRAM CONTROL, TO ANY LOCAL STORE REGISTER. CONNECTING A BUS TO A REGISTER ("SETTING A BUS CONTROL") IS A PRIMITIVE OPERATION FOR THE NANOPROGRAMMER. ONCE A BUS CONTROL HAS BEEN SET, THE BUS REMAINS CONNECTED TO THE REGISTER UNTIL THE NANOPROGRAM CHANGES THAT BUS CONTROL. THERE IS NO RESTRICTION ON THE NUMBER OR IDENTITY OF BUSES THAT MAY BE CONNECTED TO ANY (ONE) LOCAL STORE REGISTER AT THE SAME TIME, ALTHOUGH EACH BUS IS CONNECTED TO ONE AND ONLY ONE REGISTER AT ANY GIVEN TIME. ONCE A WORD APPEARS ON A BUS, HOWEVER, IT REMAINS AVAILABLE UNTIL SOME SPECIFIC ACTION CHANGES THE BUS SOURCE. (THUS, FOR EXAMPLE, IT IS POSSIBLE TO GATE THE CONTENTS OF A CONTROL STORE LOCATION INTO SEVERAL LOCAL STORE REGISTERS BY SUCCESSIVELY CHANGING THE COD BUS CONTROL AND EXECUTING THE APPROPRIATE GATE NANOPRIMITIVE.)

IF THE DATA ON TWO OR MORE BUSES ARE GATED INTO THE SAME REGISTER SIMULTANEOUSLY, THE LOGICAL "OR" OF THE VALUES APPEARS IN THE REGISTER.

A CONVENIENT MODEL OF THE BUSING STRUCTURE REPRESENTS EACH BUS CONTROL AS A "ROTARY SWITCH" ATTACHED TO A DATA PATH; THE POSITION OF THE SWITCH AS LAST SET CONNECTS THE PATH TO ONE OF THE 32 LOCAL STORE REGISTERS. A "GATE" NANOPRIMITIVE ACTIVATES DATA TRANSMISSION ON ANY PATH INTO LOCAL STORE.

SINCE THE NANOPROGRAMMER WILL TYPICALLY USE MANY OF THE LOCAL STORE REGISTERS TO SUPPORT THE FUNCTIONS OF SOME HIGHER LEVEL EMULATED MACHINE (ACCUMULATORS, LOCATION COUNTERS, MEMORY ADDRESS REGISTERS, STACK POINTERS, GENERAL-PURPOSE REGISTERS, ETC.), THE BUS CONTROLS EFFECTIVELY ALLOW THE RESOURCE ORGANIZATION OF THAT MACHINE TO BE DYNAMICALLY REDEFINED TO BEST FIT THE CURRENT TASK.

4.2.2 LOCAL STORE SPECIAL FEATURES

THE LOCAL STORE REGISTERS ARE LABELED R0 THROUGH R31. IN ADDITION TO THEIR STANDARD PROPERTIES AS MEMBERS OF LOCAL STORE, CERTAIN OF THE LOCAL STORE REGISTERS POSSESS SPECIAL CAPABILITIES AS ILLUSTRATED IN FIGURE 4.4.2A.

SPECIAL FEATURES OF LOCAL STORE

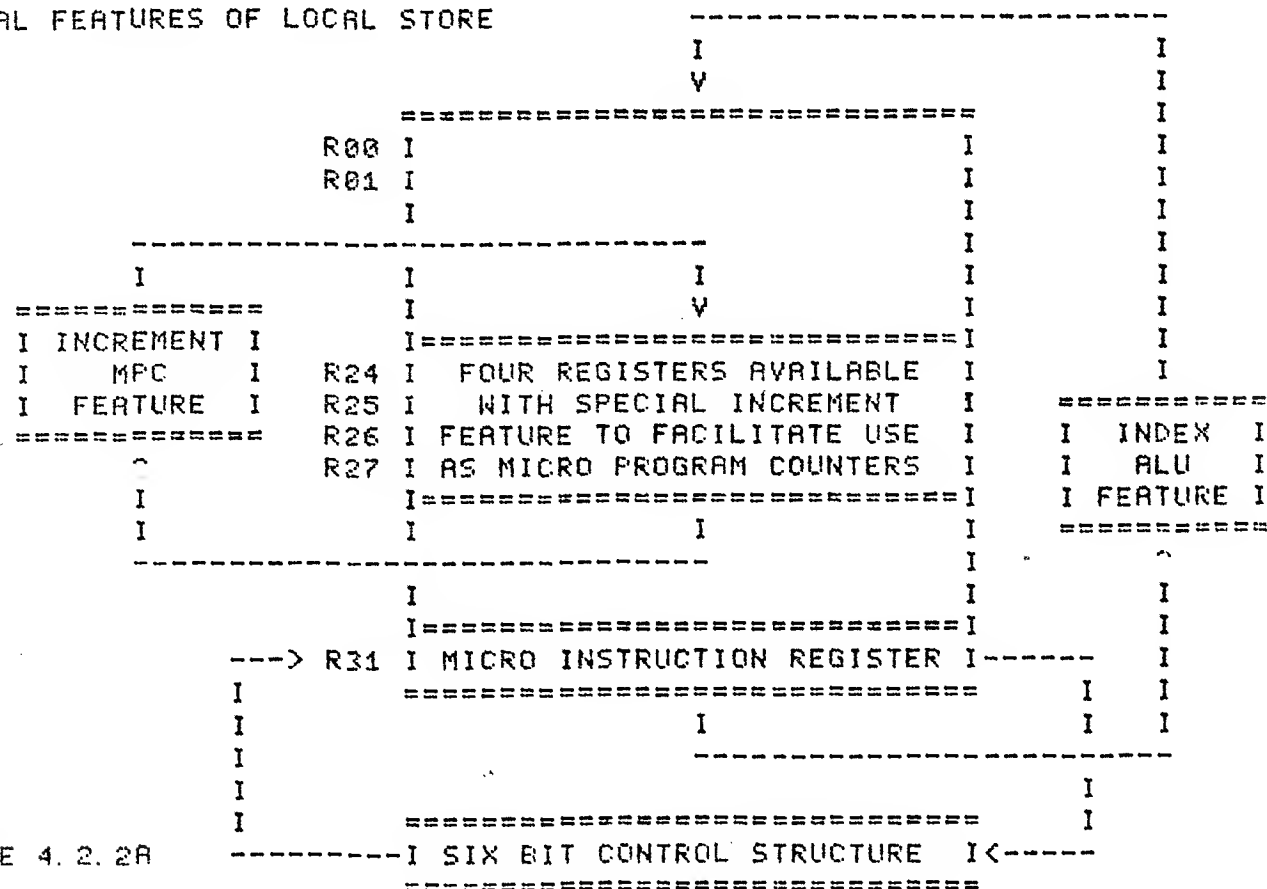


FIGURE 4.2.2A

4.2.2.1 MICRO INSTRUCTION REGISTER

THE MOST IMPORTANT SPECIAL FACILITY IN LOCAL STORE INVOLVES R31. THIS IS THE ONLY LOCAL STORE REGISTER THAT IS DEDICATED TO A SPECIFIC PURPOSE. R31 IS THE MICRO INSTRUCTION REGISTER. WHEN A CONTROL STORE WORD IS EXECUTED AS A MICROINSTRUCTION, IT CAN READILY BE GATED INTO R31 SO THAT THE NANOPROGRAMMER

MAY CONVENIENTLY MAKE USE OF THE PARAMETER INFORMATION IN THE WORD (THE MICRO-OPCODE IS AUTOMATICALLY CLEARED TO ZEROS). THUS, R31 SERVES AS THE MICRO INSTRUCTION REGISTER (MIR).

TO ALLOW MICROINSTRUCTION PARAMETERS ACCESS TO THE QM-1 SIX-BIT CONTROL STRUCTURES (TO BE PRESENTED IN SECTION 4.3), R31 IS PARTITIONED INTO THREE 6-BIT FIELDS: C, A, B (HIGH TO LOW ORDER). HENCE R31 ALSO SERVES THE SPECIAL FUNCTION OF INTERFACING THE QM-1 18-BIT AND SIX-BIT ARCHITECTURES (SEE SECTION 5.3.4).

4.2.2.2 MICRO PROGRAM COUNTERS

STILL A DIFFERENT SPECIAL CAPABILITY APPLIES TO FOUR OTHER REGISTERS IN LOCAL STORE: R24, R25, R26 AND R27. AN INCREMENT MPC FEATURE IS PROVIDED TO FACILITATE THE USE OF ANY OF THESE REGISTERS AS A "MICRO-PROGRAM COUNTER" (LOCATION COUNTER FOR MICROPROGRAMS EXECUTING OUT OF CONTROL STORE). CONTROLS EXIST FOR DIRECTLY INCREMENTING ANY OF THESE REGISTERS BY ONE OF THE FOLLOWING VALUES:

+1

+2

"B" FIELD OF R31 (SIGN-EXTENDED; 2'S COMPLEMENT)

"AB" FIELD OF R31 (LOW ORDER 11 BITS ONLY, SIGN EXTENDED)

NOTE: THE INCREMENT MPC FACILITY IS A 16 BIT ARITHMETIC UNIT; THE HIGH ORDER TWO BITS OF THE LOCAL STORE MPC REGISTER ARE SIGN EXTENDED FROM THE 16TH BIT.

OTHER ELEMENTS OF THE MPC FACILITY ARE DISCUSSED IN SECTIONS 4.2.4, 4.5.3, AND 5.6.5.

4.2.2.3 INDEX ALU FEATURE

AN "INDEX ALU" CAPABILITY IS AVAILABLE FOR ALL LOCAL STORE REGISTERS OTHER THAN THE FOUR MPC REGISTERS. ARITHMETIC OPERATIONS MAY BE PERFORMED ON THE CONTENTS OF THESE REGISTERS DIRECTLY USING ONE OF SEVERAL QUANTITIES (IN 2'S COMPLEMENT FORM) WITHOUT ROUTING THROUGH THE ARITHMETIC-LOGIC UNIT. SELECTION OF INDEX SOURCE IS MADE FROM THE FOLLOWING LIST:

ONE OF 14 EXTERNAL STORE REGISTERS
DATA ON THE CDD BUS
DATA ON THE MOD BUS

FURTHER DETAIL ON THE INDEX ALU FEATURE IS GIVEN IN SECTION 5.6.4

4.2.3 ALU AND SHIFTER

4.2.3.1 ALU

THE ARITHMETIC-LOGIC UNIT CAN BE CONTROLLED TO PERFORM ALL OF THE 16 LOGICAL (BOOLEAN) OPERATIONS, AS WELL AS CERTAIN ARITHMETIC OPERATIONS (INCLUDING ADDITION AND SUBTRACTION), UPON THE TWO 18-BIT OPERANDS PRESENT ON THE AIL AND AIR BUSES. (THE CARRY-IN VALUE FOR ARITHMETIC FUNCTIONS IS SUPPLIED BY THE CIH BIT; SEE SECTION 4.2.3.4). THE 18-BIT RESULT PROCEEDS THROUGH THE SHIFTER EXTENSION TO THE ADD BUS, WHERE IT IS AVAILABLE FOR GATING INTO A LOCAL STORE REGISTER UPON EXECUTION OF THE NANOPRIMITIVE "GATE ALU". THE ALU MAY BE USED TO DO 2'S COMPLEMENT, 1'S COMPLEMENT, OR UNSIGNED ARITHMETIC. (2'S COMPLEMENT ARITHMETIC IS MOST CONSISTENT WITH OTHER CPU MECHANISMS).

A 16 BIT MODE PERMITS THE INPUTS TO BE SIGN EXTENDED FROM 16 TO 18 BITS SO THAT THE OPERATION OF THE ALU NEED NOT BE CHANGED WHEN DEALING WITH 16 BIT DATA VALUES.

A DECIMAL CONTROL FACILITATES DECIMAL ARITHMETIC BY GENERATING A "DECIMAL CORRECTION WORD" ON THE SOD BUS WHILE BINARY FUNCTIONS ARE PERFORMED IN THE ALU. IF THE ALU PROPAGATES A CARRY OUT OF A FOUR-BIT GROUP (COUNTING FROM THE LOW-ORDER END), "0000" IS FORCED ONTO THE CORRESPONDING GROUP ON THE SOD BUS. IF NO CARRY IS PROPAGATED, "0110" IS FORCED. THE HIGH-ORDER TWO BITS OF SOD ARE FORCED TO ZEROS. WHEN THE DECIMAL CONTROL IS ACTIVE, THE SHIFTER EXTENSION IS AUTOMATICALLY BYPASSED. THE SHIFTER INPUT IS ALSO BLOCKED, AND HAS NO EFFECT ON THE CORRECTION WORD VALUE.

ALU FUNCTIONS INCLUDE PASS LEFT, FOR TRANSFERRING THE VALUE ON THE AIL BUS DIRECTLY TO THE SHIFTER EXTENSION WITHOUT INCURRING ALU PROPAGATION DELAY.

4.2.3.2 SINGLE SHIFTS

THE SHIFTER CAN BE FUNCTIONED TO PERFORM A LARGE NUMBER OF DIFFERENT SHIFT OPERATIONS UPON THE DATA PRESENT ON THE SID BUS. THE RESULT IS PLACED ON THE SOD BUS, WHERE IT IS AVAILABLE FOR GATING INTO A LOCAL STORE REGISTER UPON EXECUTION OF THE NANOPRIMITIVE "GATE SH". WHEN NO SHIFT OPERATION IS SPECIFIED, THE SHIFTER FUNCTIONS AS A DIRECT CONNECTION FROM THE SID BUS TO THE SOD BUS, PROVIDING A CONVENIENT ROUTE FOR TRANSFERS BETWEEN LOCAL STORE REGISTERS.

SHIFT OPERATIONS AS DESCRIBED ABOVE, INVOLVING ONLY THE SHIFTER AND THE SID AND SOD BUSES, ARE KNOWN AS "SINGLE-LENGTH" SHIFTS. THE FOLLOWING TYPES OF SINGLE-LENGTH SHIFTS CAN BE SPECIFIED:

LEFT AND RIGHT LOGICAL:

ZEROS INSERTED AT ONE END. BITS SHIFTED OFF THE OTHER END.

RIGHT ARITHMETIC:

SIGN BIT (HIGH-ORDER BIT) EXTENDED (COPIED) RIGHTWARD, BITS SHIFTED OFF RIGHT END.

LEFT AND RIGHT CIRCULAR:

ROTATIONS OF THE 18-BIT QUANTITY.

FOR EACH TYPE OF SHIFT OPERATION, SHIFTS OF ANY (MEANINGFUL) NUMBER OF POSITIONS ARE PERFORMED IN PARALLEL -- I.E., AS A SINGLE HARDWARE OPERATION. HENCE SINGLE-LENGTH SHIFTS OF 0 THROUGH 18 POSITIONS CAN BE PERFORMED DIRECTLY.

NOTE: WHEN SINGLE-LENGTH SHIFTS ARE SPECIFIED, THE SHIFTER EXTENSION FUNCTIONS AS A DIRECT CONNECTION FROM THE ALU OUTPUT TO THE ADD BUS.

4.2.3.3 DOUBLE SHIFTS

WHEN A DOUBLE-LENGTH SHIFT OPERATION IS SPECIFIED, THE SHIFTER EXTENSION JOINS THE SHIFTER IN TREATING THE ALU OUTPUT AND THE VALUE ON THE SID BUS AS THE HIGH-ORDER AND LOW-ORDER HALVES, RESPECTIVELY, OF A 36-BIT QUANTITY. THE HIGH-ORDER AND LOW-ORDER HALVES OF THE SHIFTED RESULT APPEAR ON THE AOD AND SOD BUSES, RESPECTIVELY. IN SOME TYPES OF DOUBLE-LENGTH SHIFTING, A 37TH BIT, INVOLVED IN THE CARRY FUNCTION, IS ALSO USED. WHEN INCLUDED, IT IS PLACED TO THE LEFT OF THE SHIFTER EXTENSION. (CARRY WILL BE FURTHER TREATED LATER.) THE FOLLOWING TYPES OF DOUBLE-LENGTH SHIFTS CAN BE SPECIFIED:

LEFT AND RIGHT LOGICAL:

ZEROS INSERTED AT ONE END OF A 37-BIT QUANTITY
(CARRY INCLUDED), BITS SHIFTED OFF OTHER END.

RIGHT ARITHMETIC:

SIGN BIT OF 36-BIT QUANTITY (HIGH-ORDER BIT OF
ALU RESULT) EXTENDED (COPIED) RIGHTWARD; BITS SHIFTED
OFF RIGHT END OF SHIFTER. CARRY IS NOT INVOLVED.

LEFT ARITHMETIC:

SAME AS LEFT LOGICAL, EXCEPT THAT THIS OPERATION CAN
SET THE OVERFLOW CONDITION (TO BE DISCUSSED), WHEREAS
LEFT LOGICAL DOES NOT SET OVERFLOW.

LEFT AND RIGHT CIRCULAR:

ROTATIONS OF THE 36-BIT QUANTITY (CARRY NOT INVOLVED).

DOUBLE-LENGTH SHIFTS OF ANY NUMBER OF POSITIONS (0 THROUGH 37) ARE
ALSO PERFORMED IN PARALLEL.

NOTE: WHEN DOUBLE LENGTH SHIFTS ARE SPECIFIED, THE AOD BUS CONTAINS THE HIGH
ORDER PORTION OF THE SHIFTED QUANTITY FOR AS LONG AS THE DOUBLE-SHIFT
IS IN CONTROL AND THE INPUTS ARE STABLE.

4.2.3.4 CARRY CONTROL

TWO FLIP-FLOPS ARE INVOLVED IN CARRY FUNCTIONS WITHIN THE ALU-SHIFTING COMPLEX: THE "CARRY-IN HOLD" (CIH) AND THE "CARRY-OUT HOLD" (COH).

THE 37TH BIT POSITION INVOLVED IN DOUBLE LOGICAL AND ARITHMETIC SHIFTS (SECTION 4.2.3.3) IS KNOWN AS THE "SH END" POSITION; IT IS LOGICALLY LOCATED AT THE HIGH-ORDER (LEFT) END OF THE DOUBLE-LENGTH SHIFT UNIT.

TWO OTHER ELEMENTS ARE REQUIRED IN THE MODEL TO BE EXPLAINED BELOW:

- A) TWO INDEPENDENTLY PROGRAMMABLE CONTROLS, "LEFT CTL" AND "RIGHT CTL"
- B) THE FOLLOWING MUTUALLY EXCLUSIVE NANOPRIMITIVE OPERATIONS:

SET CIH	CLEAR CIH	SH TO COH
SET COH	CLEAR COH	ALU TO COH
		ALU TO BOTH (COH AND CIH)

FIGURE 4.2.3.4A AIDS IN UNDERSTANDING THE ALU-SHIFT-CARRY SYSTEM.

THE OUTPUT OF CIH IS PERMANENTLY ENABLED AS ALU CARRY-IN, AND HAS NO OTHER FUNCTION.

CIH CAN BE LOADED FROM ONE OF TWO SOURCES:

- A) DIRECT PROGRAM LOAD: "SET CIH", "CLEAR CIH".
- B) ALU CARRY-OUT; EFFECTED BY "ALU TO BOTH".

THE OUTPUT OF COH IS PERMANENTLY ENABLED TO SERVE THE FOLLOWING FUNCTIONS:

- A) SOLE INPUT TO THE SH END BIT POSITION.
- B) ONE OF TWO INPUTS TO THE "LEFT CONTROL SWITCH", TO BE EXPLAINED.
- C) SOLE SOURCE OF THE "CARRY TEST" VALUE, ONE OF THE "LOCAL CONDITIONS" THAT CAN BE TESTED IN A NANOPROGRAM.

COH CAN BE LOADED FROM ONE OF THREE SOURCES:

- A) DIRECT NANOPROGRAM LOAD: "SET COH", "CLEAR COH".
- B) ALU CARRY-OUT; EFFECTED BY "ALU TO COH" OR "ALU TO BOTH".
- C) OUTPUT OF THE "RIGHT CONTROL SWITCH", TO BE EXPLAINED.

THE OUTPUT OF THE LEFT CONTROL SWITCH IS THE LOW-ORDER BIT POSITION OF THE SOD BUS. IN ITS NORMAL STATE, THIS SWITCH CONNECTS THE SOD-LOW-BIT TO THE LOW-ORDER BIT POSITION OF THE SHIFTER OUTPUT. WHEN THE "LEFT CTL" IS ACTIVE, HOWEVER, THE LATTER CONNECTION IS BROKEN, AND SOD-LOW-BIT IS INSTEAD TAKEN

4.2.3.5 TEST CONDITIONS

INCLUDING CARRY, THERE ARE SIX "LOCAL CONDITIONS", GENERATED BY ALU AND/OR SHIFTING OPERATIONS, WHICH CAN BE TESTED BY NANOPRIMITIVES. THEY ARE:

CARRY (C) OUTPUT OF COH, AS DISCUSSED IN SECTION 4.2.3.4.

SIGN (S) THE HIGH-ORDER BIT ON THE AOD BUS.

RESULT (R) NORMALLY THE LOGICAL OR OF THE LOW-ORDER 17 BITS ON THE AOD BUS. HOWEVER, WHEN EITHER OF THE SPECIAL CARRY CONTROLS, "RIGHT CTL" OR "LEFT CTL", ARE SET, "RESULT" IS THE LOGICAL OR OF THE LOW ORDER 17 BITS ON THE AOD BUS AND ALL 18 BITS ON THE SOD BUS. THUS AN ABSOLUTE ZERO VALUE, EITHER 18 OR 36 BITS MAY BE TESTED WITH THE CONDITION OF BOTH S AND R EQUAL 0.

OVERFLOW (O) THIS CONDITION IS THE LOGICAL OR OF SHIFTING OVERFLOW AND ALU OVERFLOW. SHIFTING OVERFLOW ARISES ONLY IN DOUBLE LEFT ARITHMETIC SHIFTS, AND IS DEFINED TO ARISE IF AND ONLY IF A SERIAL (BIT-BY-BIT) SHIFT OF THE SAME NUMBER OF POSITIONS WOULD, AT ANY TIME, CHANGE THE VALUE OF THE HIGH-ORDER (SIGN) BIT OF AOD. ALU OVERFLOW ARISES (SEE SECTION 5.6.2), IF AND ONLY IF THE BIT-CARRY SIGNALS PROPAGATED INTO THE SIGN AND CARRY-OUT POSITIONS ARE OF OPPOSITE VALUES.

SHIFTER HIGH BIT (SHB) THE HIGH-ORDER BIT ON THE SOD BUS.

SHIFTER LOW BIT (SLB) THE LOW-ORDER BIT ON THE SOD BUS.

SINCE IT IS HIGHLY DESIRABLE TO HAVE A CONVENIENT METHOD OF PRESERVING CONDITION BITS, A LOCATION IS PROVIDED IN THE CPU FOR A COPY OF EACH OF THE SIX LOCAL CONDITIONS. THE NANOPROGRAMMER CAN SET A CONTROL SUCH THAT WHEN GATE ALU IS EXECUTED, THE FOUR LOCAL CONDITIONS C, S, R, AND O ARE AUTOMATICALLY COPIED INTO THEIR CORRESPONDING "GLOBAL CONDITION" BITS, AND (UNDER SEPARATE CONTROL) WHEN GATE SH IS EXECUTED, THE TWO LOCAL CONDITIONS SHB AND SLB ARE COPIED INTO THEIR GLOBAL COUNTERPARTS. THEN, INDEPENDENT NANOPRIMITIVE TESTS CAN BE MADE UPON THESE "GLOBAL CONDITIONS".

4.2.3.6 SIXTEEN-BIT MODE

A SPECIAL CPU FEATURE IS INCLUDED TO FACILITATE MANIPULATION OF BYTE-ORIENTED DATA. A "16-BIT MODE" CONTROL CAN BE SET BY THE NANOPROGRAMMER, WITH THE FOLLOWING EFFECTS:

- A) THE LOCAL CONDITIONS S, R, AND SHB ARE REDEFINED TO FUNCTION AS IF THE SHIFTER AND SHIFTER EXTENSION WERE EACH 16 BITS WIDE, WITH THE VIRTUAL UNITS MAPPED ONTO THE LOW-ORDER 16 BITS OF THE 18-BIT UNITS (I.E., THE S-TEST IS TAKEN FROM ADD BIT 15 INSTEAD OF ADD BIT 17 -- USING 0-ORIGIN NUMBERING; ETC.). THE "DOUBLE WIDTH" R TEST IS BASED ON THE LOWER 32 BITS OF THE CONCATENATED ADD AND SOD BUSES.
- B) THE RIGHT CONTROL SWITCH SELECTS BIT 16 OF THE ADD BUS INSTEAD OF THE OUTPUT OF THE SH END BIT TO LOAD COH.
- C) THE ALU OVERFLOW CONDITION AND ALU CARRY-OUT ARE REDEFINED TO FUNCTION AS IF THE ALU WERE 16 BITS WIDE, WITH THE VIRTUAL ALU MAPPED ONTO THE LOW-ORDER 16 BITS OF THE 18-BIT ALU.
NOTE: SHIFTING OVERFLOW IS NOT REDEFINED.
- D) THE AIL AND AIR BUSES AUTOMATICALLY COPY THE 3RD-HIGHEST-ORDER BIT (I.E., THE SIGN BIT OF A 16-BIT WORD) INTO THE TWO HIGH-ORDER POSITIONS; THUS ARITHMETIC IN 16-BIT MODE ALSO GENERATES CORRECT 18-BIT RESULTS, FOR LATER USE IN 18-BIT MODE IF DESIRED.
IMPORTANT: THE "PASS LEFT" ALU FUNCTION, WHICH ROUTES THE CONTENTS OF AIL DIRECTLY AROUND THE ALU, ALSO BYPASSES THIS SIGN-EXTENSION MECHANISM.

4.2.4 CONTROL STORE

CONTROL STORE IS A FULLY READABLE/WRITABLE GENERAL-PURPOSE 18-BIT WIDE STORE, IMPLEMENTED IN SEMICONDUCTOR MEMORY. IT IS AVAILABLE IN BLOCKS OF 1K WORDS, UP TO A MAXIMUM OF 16K WORDS.

THE NANOPRIMITIVES "READ CS", "WRITE CS", AND "GATE CS" ARE PROVIDED TO ACCESS CONTROL STORE. THE READ CS AND WRITE CS NANOPRIMITIVES ARE ACCOMPANIED BY A FIELD WHICH SELECTS THE SOURCE OF THE ADDRESS IN CONTROL STORE AT WHICH A WORD IS TO BE ACCESSED. SOURCES OF CS ADDRESSES ARE AS FOLLOW:

- CIA: THE VALUE ON THE CIA BUS; FOR GENERAL-PURPOSE DATA ACCESS.
- COD: THE VALUE ON THE COD BUS; FOR CONVENIENT INDIRECT ACCESS.
- MPC, MPC+1, MPC+2, MPC+B, MPC+AB (LOW-ORDER 11 BITS OF R31):
(INCREMENTS SIGN-EXTENDED, 2'S COMPLEMENT.)
FOR MICROINSTRUCTION SEQUENCING AND BRANCHING,
AND FOR READING MICROINSTRUCTION PARAMETER LISTS;
MICROINSTRUCTION EXECUTION IS DISCUSSED IN
SECTION 4.5.3. SELECTION OF WHICH OF THE
FOUR MPC'S IS TO BE USED IS MADE BY A
MECHANISM SIMILAR TO A BUS CONTROL (SEE SECTION 4.3.2.3)
- SWITCHES: THE (12-BIT) VALUE TAKEN FROM MANUAL SWITCHES
IN THE MACHINE CABINET.

WHEN A WORD HAS BEEN READ OUT OF CONTROL STORE, IT APPEARS ON THE COD BUS, AVAILABLE FOR GATING INTO A LOCAL STORE REGISTER BY EXECUTION OF THE NANOPRIMITIVE GATE CS. ONCE ESTABLISHED, A COD VALUE REMAINS UNTIL CHANGED BY THE NEXT READ CS OPERATION.

WRITING A WORD INTO CONTROL STORE IS ACCOMPLISHED BY PLACING THE DATUM ON THE CID BUS, AND THEN EXECUTING THE NANOPRIMITIVE WRITE CS WITH THE APPROPRIATE CS ADDRESS SELECTED.

IF READ CS AND WRITE CS ARE EXECUTED SIMULTANEOUSLY, READ CS IS IGNORED.

EXECUTION OF READ CS FROM A NONEXISTENT LOCATION GENERATES ZEROS ON THE COD BUS; EXECUTION OF WRITE CS AT A NONEXISTENT LOCATION IS A NULL OPERATION. DUE TO THE MAXIMUM ALLOWABLE ADDRESS SPACE FOR CONTROL STORE, HOWEVER, ALL CS-ADDRESSES ARE TREATED MODULO 32K (I.E., THE HIGH-ORDER THREE BITS ARE IGNORED).

4.2.5 EXTERNAL STORE

EXTERNAL STORE IS A BANK OF 32 REGISTERS, PARTITIONED INTO SEVERAL GROUPS TO SUPPORT SPECIFIC FUNCTIONS: EXTERNAL PORTS, INDEX REGISTERS, MAIN STORE ADDRESSING FACILITIES, AND INTERRUPT CONTROL.

ALTHOUGH EACH TYPE OF ES REGISTER IS ASSOCIATED WITH SPECIAL HARDWARE FACILITIES TO IMPLEMENT ITS SPECIFIC FUNCTION, ALL 32 ES REGISTERS ARE UNIFORMLY ACCESSIBLE BY THE NANOPROGRAM VIA THE EOD AND EID BUSES. TO PROVIDE THIS ACCESSIBILITY, TWO ADDITIONAL BUS CONTROLS ARE ASSOCIATED WITH EXTERNAL STORE TRANSFERS, AS FOLLOWS:

WHILE THE DESTINATION END OF THE EOD BUS IS CONNECTED TO ONE OF THE 32 LOCAL STORE REGISTERS BY THE EOD BUS CONTROL, THE SOURCE END OF THE SAME BUS IS CONNECTED TO ONE OF THE 32 EXTERNAL STORE REGISTERS BY A DIFFERENT BUS CONTROL, LABELED EOA. THE TRANSFER FROM ES TO LS IS EXECUTED BY THE NANOPRIMITIVE "GATE ES". SIMILARLY, THE EXTERNAL STORE CONNECTION OF THE EID BUS IS SELECTED BY THE EIA BUS CONTROL. THE TRANSFER FROM LS TO ES IS EXECUTED BY THE NANOPRIMITIVE "LOAD ES".

FUNCTIONS OF EXTERNAL STORE REGISTERS ARE AS FOLLOWS:

- E0 THROUGH E7: EIGHT PORT REGISTERS AVAILABLE FOR INTERFACING THE QM-1 TO ITS ENVIRONMENT. THESE REGISTERS ARE DIRECTLY CONNECTABLE TO MAIN STORE. (THE QM-1 EXTERNAL INTERFACE IS DISCUSSED IN SECTION 4.6.)
- E8 THROUGH E21: FOURTEEN INDEX ALU OPERAND SOURCES. THESE INCLUDE EIGHT REGISTERS FOR GENERAL USE AND SIX REGISTERS FROM THE GROUPS BELOW.
- E16, E17: BASE ADDRESS REGISTER AND FIELD LENGTH REGISTER ASSOCIATED WITH MAIN STORE ADDRESSING (DISCUSSED IN SECTION 4.2.6.3). INCLUSION OF THESE MACHINE-CONTROL FUNCTIONS IN E16 AND E17 IS A QM-1 OPTION; IF SUCH FUNCTIONS ARE NOT INCLUDED, E16 AND E17 ARE SCRATCH REGISTERS.
- E18 THROUGH E31: FOURTEEN REGISTERS ASSOCIATED WITH INTERRUPT CONTROL. INTERRUPT FACILITIES ARE DISCUSSED IN SECTION 4.5.2.4. THE REGISTERS INCLUDE:
 - E18 AND E19: INTERRUPT ENABLE BITS.
 - E20 AND E21: INTERRUPT PENDING FLAGS.
 - E22 --- E31: INTERRUPT ADDRESS FIELDS.

THE OVERALL LAYOUT OF EXTERNAL STORE IS SHOWN IN FIGURE 4.2.5A

	I	-----	I	-----	
E0	I		I		R
	I	-----	I		E
E1	I		I		P
	I	-----	I		G
					O
					R
					S
	I	-----	I		T
E6	I		I		T
	I	-----	I		E
E7	I		I		R
	I	-----	I		S
E8	I		I		
	I	-----	I		I
					INDEX
					ALU
					OPERANDS
E15	I		I		I
	I	-----	I		I
E16	I	MAIN STORE BASE ADDRESS REGISTER	I	MAIN	I
	I	-----	I	STORE	I
E17	I	MAIN STORE FIELD LENGTH REGISTER	I	CONTROLS	I
	I	-----	I	-----	I
E18	I	PROGRAM/02 03 04/05 06 07/08 09 10/11 12 13/14 15 16	I	INTERRUPT	I
	I	CHECK /-----	I	ENABLE	I
E19	I	MASKS /17 18 19/20 21 22/23 24 25/26 27 28/29 30 31	I	CONTROLS	V
	I	-----	I	-----	
E20	I	PROGRAM/02 03 04/05 06 07/08 09 10/11 12 13/14 15 16	I	INTERRUPT	
	I	CHECK /-----	I	PENDING	
E21	I	FLAGS /17 18 19/20 21 22/23 24 25/26 27 28/29 30 31	I	FLAGS	
	I	-----	I	-----	
E22	I	02 / 03 / 04	I	I	A
	I	-----	I	N	D
E23	I	05 / 06 / 07	I	T	D
	I	-----	I	E	R
			I	R	E
	I	-----	I	R	S
E30	I	26 / 27 / 28	I	U	S
	I	-----	I	P	E
E31	I	29 / 30 / 31	I	T	S
	I	-----	I	-----	
BIT	17 16 15 14 13 12 11 10 09 08 07 06 05 04 03 02 01 00				

LAYOUT OF EXTERNAL STORE

FIGURE 4.2.5A

4.2.6 MAIN STORE

4.2.6.1 GENERAL

MAIN STORE IS A GENERAL-PURPOSE 18-BIT-WIDE CORE STORAGE, AVAILABLE IN BLOCKS OF 8K WORDS UP TO 256K MAXIMUM (16K WORDS MINIMUM). FULL CYCLE TIME IS 800 NANoseconds; SINCE LOWER-LEVEL CONTROL OPERATIONS OCCUR AN ORDER OF MAGNITUDE FASTER, MAIN STORE IS WELL SUITED TO CONTAIN PROGRAMS OF VIRTUAL (EMULATED) MACHINES WHOSE INSTRUCTIONS REQUIRE A MODERATELY COMPLICATED INTERPRETATION AT LOWER LEVELS.

FOR CONVENIENCE AND EFFICIENCY IN INPUT/OUTPUT PROCESSING, THE TWO BUSES ASSOCIATED WITH MAIN STORE (MIX AND MOD) MAY BE CONNECTED NOT ONLY TO ANY OF THE 32 LOCAL STORE REGISTERS, BUT ALSO TO ANY OF THE 8 PORT REGISTERS IN EXTERNAL STORE. THUS FOR THE MIX AND MOD BUS CONTROLS ONLY, THE PORT REGISTERS ARE TREATED AS EXTENSIONS TO LOCAL STORE; THEY ARE DESIGNATED AS R32 THROUGH R39 WHEN USED IN THIS WAY.

4.2.6.2 MS OPERATIONS

TO INITIATE A FULL (NON-DESTRUCTIVE) READ OPERATION IN MAIN STORE, THE NANOPROGRAMMER FIRST DETERMINES THAT MAIN STORE IS NOT BUSY ("MS BUSY" IS ONE OF THE "SPECIAL CONDITIONS" AVAILABLE FOR NANOPROGRAM TESTING), AND THEN SIMULTANEOUSLY EXECUTES THE TWO NANOPRIMITIVES "MSGO" AND "MSRS" (FOR "MAIN STORE RESTART"). MAIN STORE ACCESSES THE LOCATION ADDRESSED BY THE VALUE ON THE MIX BUS, AS MODIFIED BY ADDRESSING FACILITIES WHICH ARE DISCUSSED IN THE NEXT SECTION.

WHEN THE ACCESSED WORD IS AVAILABLE, ANOTHER SPECIAL TEST CONDITION, "MS DATA INVALID", BECOMES FALSE, AND THE NANOPROGRAM CAN GATE OUT THE WORD THROUGH THE MOD BUS BY EXECUTING THE NANOPRIMITIVE "GATE MS"; ACCESS TIME IS 640 NANoseconds.

THE SAME TWO NANOPRIMITIVES, MSGO AND MSRS, ARE USED TO CONTROL OTHER OPERATIONS OF MAIN STORE, AS FOLLOWS:

WHEN MSGO IS EXECUTED WITHOUT MSRS, MAIN STORE BEGINS THE FIRST HALF-CYCLE ("EXTRACTION PART") OF A SPLIT-CYCLE OPERATION. AS IN THE CASE OF A FULL-READ OPERATION, THE ADDRESS IS TAKEN FROM THE MIX BUS. MAIN STORE WILL REMAIN BUSY UNTIL THE COMPLETION OF THE SECOND HALF-CYCLE ("INSERTION PART"). IN THE SPLIT-CYCLE MODE OF OPERATION, HOWEVER, THE LATTER MUST BE EXPLICITLY INVOKED BY THE NANOPROGRAM EXECUTION OF MSRS; THE DATA WORD TO BE INSERTED (WRITTEN) IS TAKEN FROM THE MIX BUS AT THE TIME MSRS IS EXECUTED. THIS MODE OF OPERATION MAY BE USED IN TWO WAYS:

TO PERFORM A READ/MODIFY/WRITE SEQUENCE, THE NANOPROGRAMMER INITIATES SPLIT-CYCLE OPERATION AS DESCRIBED ABOVE AND THEN, WHEN MS DATA INVALID BECOMES FALSE, GATES OUT THE EXTRACTED WORD (GATE MS) FOR MODIFICATION (FOR EXAMPLE INDEXING) BY CPU FACILITIES. WHEN THE MODIFIED WORD IS READY FOR INSERTION BACK INTO ITS MS LOCATION, IT IS PLACED ON THE MIX BUS AND MSRS IS EXECUTED. SINCE MS BUSY WILL BECOME FALSE AT THE COMPLETION OF THE SECOND HALF-CYCLE, THIS OPERATION CAN OFFER SIGNIFICANT TIME SAVINGS OVER THE ALTERNATIVE FULL-READ, DATA MANIPULATION, FULL-WRITE SEQUENCE, ESPECIALLY IF THE MODIFICATION PERIOD IS RELATIVELY SHORT.

ALTERNATIVELY, TO PERFORM A "FULL WRITE", THE NANOPROGRAMMER INITIATES SPLIT CYCLE OPERATION AS DESCRIBED ABOVE AND THEN AT ANY TIME THEREAFTER -- INCLUDING IMMEDIATELY AFTER MSGO -- PLACES THE WORD TO BE INSERTED ON THE MIX BUS AND EXECUTES MSRS. IF MSRS IS RECEIVED DURING THE FIRST HALF-CYCLE OF SPLIT-CYCLE OPERATION, MAIN STORE WILL LATCH THE DATA-IN WORD FROM

THE MIX BUS AND "REMEMBER" TO INITIATE THE SECOND HALF-CYCLE AS SOON AS POSSIBLE.

A DISTINCT ADVANTAGE TO LESSENING THE DISTINCTION BETWEEN THE FULL-WRITE AND READ/MODIFY/WRITE FUNCTIONS IS THAT THE NANOPROGRAM CAN INITIATE A MAIN STORE OPERATION WITHOUT MAKING A COMMITMENT TO EITHER OF THE TWO FUNCTIONS; IF THE NANOPROGRAM SUBSEQUENTLY DECIDES THAT THE OPERATION IS TO BE A FULL-WRITE, NO TIME LOSS IS INCURRED IF THE DECISION IS MADE BEFORE THE END OF THE FIRST HALF-CYCLE. (IN FACT, THE FULL-READ FUNCTION MAY ALSO BE ACHIEVED IN THE SPLIT-CYCLE MODE, ALTHOUGH WITH A SLIGHT DEGRADATION IN CYCLE TIME DUE TO ROUTING DELAYS.) THIS FACILITY IS THUS USEFUL IN IMPLEMENTING CERTAIN MAIN STORE MODIFICATION LOOK-AHEAD SCHEMES.

NOTES:

1. MAIN STORE IGNORES ANY MSGO SIGNAL RECEIVED WHEN MS BUSY IS TRUE.
(SEE SECTION 5.4.3.)
2. MAIN STORE IGNORES ANY MSRS SIGNAL RECEIVED WHEN EITHER:
 - A) MS IS NOT BUSY; OR
 - B) MS IS BUSY IN FULL READ MODE; OR
 - C) THE SECOND HALF-CYCLE IN SPLIT-CYCLE MODE IS ALREADY IN PROGRESS. (SEE SECTION 5.4.3.)
3. WHEN EITHER MODE OF OPERATION IS INITIATED:
 - A) MS BUSY BECOMES TRUE;
 - B) MS DATA INVALID BECOMES TRUE; AND
 - C) MOD IS CLEARED TO ZEROS.
4. WHEN THE SECOND HALF-CYCLE OF SPLIT-CYCLE MODE IS INITIATED, MOD TAKES ON THE VALUE OF THE WORD BEING INSERTED. NOTE, HOWEVER, THAT MS DATA INVALID IS SET "TRUE".

4.2.6.3 MS ADDRESSING AND PROTECTION (QM-1 OPTION)

A QM-1 OPTION IS THE USE OF EXTERNAL STORE REGISTER 16 AS THE BASE REGISTER AND EXTERNAL STORE REGISTER 17 AS THE FIELD LENGTH REGISTER ASSOCIATED WITH THE BASE-ADDRESSING, WRITE-PROTECTION, AND ADDRESS-ALARM FACILITIES TO BE DESCRIBED BELOW. IF THIS OPTION IS NOT PRESENT, THEN:

- A) E16 AND E17 ARE AVAILABLE FOR SCRATCH USE; AND
- B) THE FACILITIES DESCRIBED BELOW OPERATE AS IF THE BASE REGISTER PERMANENTLY CONTAINED THE VALUE ZERO AND THE FIELD LENGTH REGISTER PERMANENTLY CONTAINED THE VALUE $2^{18}-1$.

WHENEVER MAIN STORE USES THE VALUE ON THE MIX BUS AS AN ADDRESS, THAT VALUE IS TREATED AS A DISPLACEMENT; IT IS ADDED TO THE CONTENTS OF THE BASE REGISTER TO YIELD THE TRUE (ABSOLUTE) ADDRESS TO BE ACCESSED.

AN MS ADDRESS VIOLATION PROGRAM CHECK IS GENERATED IN EITHER OF THE FOLLOWING TWO CASES:

- A) WHEN AN MS ACCESS OF ANY KIND USES AN ABSOLUTE ADDRESS WHICH FALLS OUTSIDE THE ALLOWED RANGE DEFINED BY THE BASE AND FIELD LENGTH REGISTERS. THE LOWEST ADDRESS IN THE ALLOWED RANGE IS THE VALUE OF THE BASE REGISTER; THE NUMBER OF WORDS (CONSECUTIVE LOCATIONS) IN THE ALLOWED RANGE IS GIVEN BY THE CONTENTS OF THE FIELD LENGTH REGISTER (SO THAT ALL ACCESSES ARE VIOLATIONS WHEN THE FIELD LENGTH REGISTER IS SET TO ZERO). WRAPAROUND OF ANY KIND IS DISALLOWED.
- B) WHEN AN MS ACCESS OF ANY KIND ADDRESSES A LOCATION WHICH IS NOT PHYSICALLY PRESENT IN MAIN STORE.

FOR THE CONVENIENCE OF PROGRAMS USED AS "PRIVILEGED" OR "SYSTEM" ROUTINES, A NANOPRIMITIVE CONTROL ("DIRECT MS ACCESS") CAN MOMENTARILY FORCE THE EFFECTIVE VALUE OF THE BASE REGISTER TO ZERO AND THE EFFECTIVE VALUE OF THE FIELD LENGTH REGISTER TO $2^{18}-1$ (E16 AND E17 THEMSELVES DO NOT CHANGE VALUE).

IN ADDITION TO GENERATING THE PROGRAM CHECK, DETECTION OF MS ADDRESS VIOLATION SETS MOD TO ALL ONES AND LEAVES THE CONTENTS OF THE MEMORY UNCHANGED.

IN ALL MODES OF MAIN STORE OPERATION, A PROGRAM CHECK IS GENERATED IN THE CASE OF FAILURE OF A PARITY TEST AUTOMATICALLY MADE UPON THE EXTRACTED WORD.

4.2.6.4 RMI UNIT (QM-1 OPTION) FURTHER SPECIFICATIONS TO BE ANNOUNCED)

IF DESIRED, THE CONTENTS OF THE MOD BUS CAN BE ROUTED THROUGH THE RMI UNIT BEFORE BEING GATED INTO A LOCAL STORE REGISTER OR A CHANNEL REGISTER.

THE RMI UNIT - ROTATE, MASK, AND INDEX - IS A GENERAL-PURPOSE DATA-TRANSFORMATION DEVICE WITH SPECIAL APPLICATION IN EXTRACTING FIELDS AND DECODING INFORMATION WHEN EMULATING A "MAIN STORE MACHINE". THE OPERATION OF THE RMI UNIT PASSES A WORD THROUGH THREE SUCCESSIVE STAGES OF TRANSFORMATION:

- A) THE INITIAL VALUE UNDERGOES A RIGHT CIRCULAR SHIFT BY THE NUMBER OF POSITIONS SPECIFIED IN A ROTATE PARAMETER;
- B) THE RESULT OF THIS OPERATION IS LOGICALLY AND'ED WITH AN 18-BIT MASK PARAMETER;
- C) THE RESULT OF THIS OPERATION IS ADDED (2'S COMPLEMENT) TO AN 18-BIT INDEX PARAMETER TO YIELD THE FINAL RESULT.

THERE ARE THREE SETS OF SUCH PARAMETERS. SELECTION OF ONE OF THE THREE PARAMETER SETS IS ASSOCIATED WITH THE GATE MS NANOPRIMITIVE; A FOURTH OPTION IS TO BYPASS THE RMI UNIT.

NOTES:

1. SINCE THE MOD SOURCE VALUE REMAINS STABLE UNTIL CHANGED BY A MAIN STORE OPERATION, THE SAME WORD MAY BE TAKEN THROUGH -- A SUCCESSION OF DIFFERENT RMI TRANSFORMATIONS (AND ALSO ROUTED TO DIFFERENT DESTINATIONS) WITHOUT RE-CYCLING MAIN STORE.
2. THE RMI PARAMETERS MAY BE CHANGED AS OFTEN AS DESIRED BY THE NANOPROGRAMMER.

4.3 SIX-BIT CONTROL STRUCTURE

4.3.1 GENERAL

THE LARGE NUMBER OF HARDWARE RESOURCES IN THE QM-1 AND THE FLEXIBILITY WITH WHICH THEY MAY BE USED REQUIRE A LARGE VARIETY OF CONTROL INFORMATION, DYNAMICALLY CHANGEABLE DURING THE EXECUTION OF USER PROGRAMS.

RATHER THAN HAVING ALL SUCH INFORMATION PLACED IN A STORE FROM WHICH INSTRUCTIONS ARE EXECUTED, THE CONCEPT OF "RESIDUAL CONTROL" IS IMPLEMENTED. REGISTERS ARE PROVIDED IN THE MACHINE FOR HOLDING THIS HARDWARE-CONTROLLING INFORMATION. THESE REGISTERS CAN BE LOADED AT THE EXPLICIT COMMAND OF AN EXECUTING PROGRAM; THEIR CONTENTS WILL REMAIN IN CONTROL OF THEIR ASSIGNED HARDWARE FUNCTIONS UNTIL RELOADED (HENCE "RESIDUAL CONTROL"). IN THIS MANUAL, THE TERMS "RESIDUAL CONTROL" AND "NANOPRIMITIVE CONTROL" ARE USED WITH MUTUALLY EXCLUSIVE DEFINITIONS.

RESIDUAL CONTROL FUNCTIONS IN THE QM-1 ARE MAINTAINED IN A BANK OF SIX-BIT REGISTERS KNOWN AS F-STORE. A COMPLETE SYSTEM OF NANOPRIMITIVE CONTROLS AND SIX-BIT DATA PATHS EXISTS FOR TRANSFERRING QUANTITIES BETWEEN F-REGISTERS AND A SET OF SIX-BIT SOURCE AND DESTINATION FIELDS ELSEWHERE IN THE MACHINE, AND FOR MANIPULATING THESE DATA.

THE SIX-BIT SOURCE AND DESTINATION FIELDS (FROM THE POINT OF VIEW OF F-STORE) ARE COLLECTIVELY KNOWN AS AUXILIARY (AUX) FIELDS, ALTHOUGH SOME ARE CONTROL REGISTERS IN THEIR OWN RIGHT.

THE SAME RULES OF SIMULTANEOUS BUSING APPLY TO BOTH 18-BIT AND SIX-BIT TRANSFERS, SUCH THAT IF TWO OR MORE SOURCES ARE GATED INTO THE SAME DESTINATION SIMULTANEOUSLY, THE LOGICAL "OR" OF THOSE SOURCE VALUES APPEARS AT THE DESTINATION.

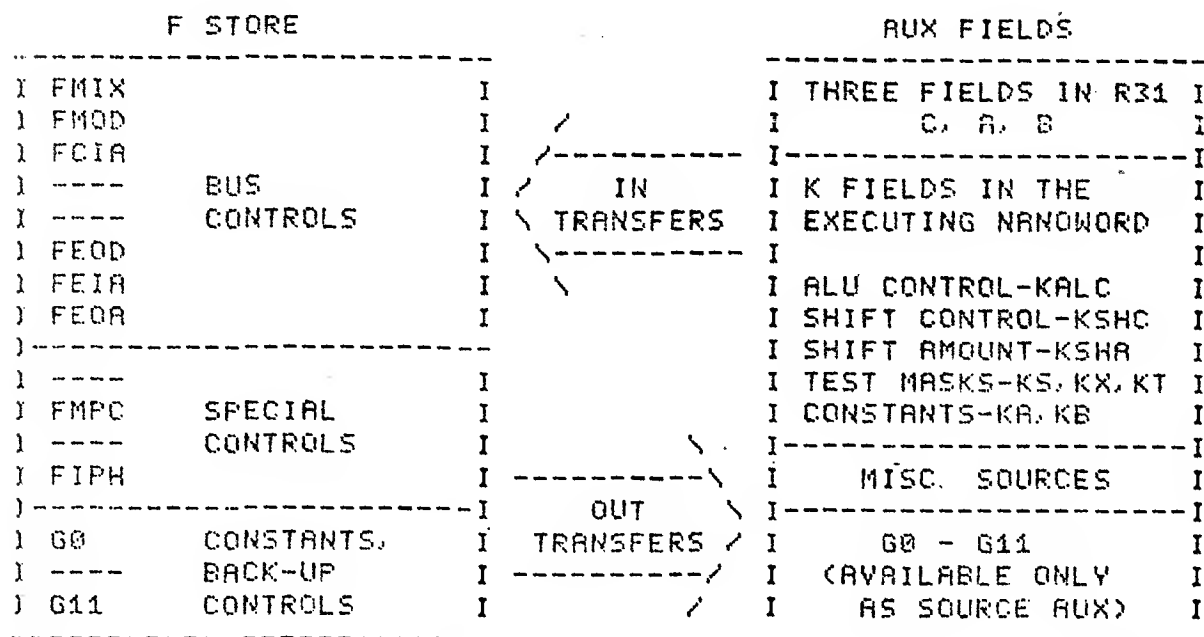
4.3.2 F-STORE

4.3.2.1 GENERAL

THE 32 SIX-BIT REGISTERS IN F-STORE, NUMBERED F0 THROUGH F31, ARE ALL UNIFORMLY ACCESSIBLE FOR THE PURPOSE OF LOADING FROM SIX-BIT SOURCE FIELDS AND READING INTO SIX-BIT DESTINATION FIELDS. EXECUTION OF SUCH SIX-BIT TRANSFERS AND THE ASSOCIATED ADDRESSING IN F-STORE (AS WELL AS SELECTION OF SOURCE AND DESTINATION FIELDS) ARE ACCOMPLISHED ENTIRELY BY NANOPRIMITIVE CONTROL.

SIMILARLY, NANOPRIMITIVE CONTROLS MAY BE APPLIED UNIFORMLY TO ANY REGISTER IN F-STORE TO INCREMENT (BY ONE) OR DECREMENT (BY ONE) THE CONTENTS OF THAT REGISTER (MODULO 64).

IT IS CONVENIENT, HOWEVER, TO APPROACH F-STORE AS PARTITIONED INTO THREE GROUPS, BY FUNCTION: BUS CONTROL F'S, SPECIAL F'S, AND G'S. THIS IS SHOWN IN FIGURE 4.3.2A.



4.3.2.2 BUS CONTROL F'S

THE FIRST 14 F-REGISTERS ARE THE BUS CONTROLS (SEE SECTION 4.2.1). THEY ARE SYMBOLICALLY REFERENCED IN ASSOCIATION WITH THEIR BUS NAMES (FMD, FAD, ETC.), WITH THE ADDITION OF FEIA AND FEOR. THE CONTENTS OF THESE REGISTERS ARE INTERPRETED IN ONE OF THREE WAYS TO ACHIEVE BUS CONTROL, DEPENDING ON THE NATURE OF THE ASSOCIATED BUS.

THE CONTENTS OF AN F-REGISTER ASSOCIATED WITH THE DESTINATION END OF A BUS (WITH THE EXCEPTION OF FMD) ARE USED MODULO 32 TO ADDRESS (CONNECT) A LOCAL STORE REGISTER (FSOD, FAD, FEOD, FOD) OR, IN THE CASE OF FEIA, AN EXTERNAL STORE REGISTER. (THE HIGH-ORDER BIT IS IGNORED FOR BUS CONTROL PURPOSES, ALTHOUGH IT IS PHYSICALLY PRESENT IN THE F-REGISTER AS LOADED.)

THE CONTENTS OF AN F-REGISTER ASSOCIATED WITH THE SOURCE END OF A BUS (WITH THE EXCEPTION OF FMIX) ARE USED MODULO 64 TO ADDRESS (CONNECT) A LOCAL STORE REGISTER (FSID, FAI, FAIR, FEID, FCIA, FCI). IF THE ADDRESS IS GREATER THAN 31 (I. E., IF THE HIGH-ORDER BIT IS SET), THE BUS IS CONNECTED TO A PERMANENT SOURCE OF ALL ONES, RATHER THAN TO A LOCAL STORE REGISTER. IN THE CASE OF FEOR GREATER THAN 31, THE EOD BUS IS CONNECTED TO A SOURCE OF ALL ZEROS RATHER THAN TO AN EXTERNAL STORE REGISTER.

SINCE MIX AND MOD HAVE AN ADDRESSING RANGE BEYOND 32 (SEE SECTION 4.2.6.1), SPECIAL RULES ARE USED IN INTERPRETING THE BUS CONTROL FUNCTIONS OF FMIX AND FMD; THESE F-REGISTERS ARE USED MODULO 64, WITH THE EIGHT PORT REGISTERS (E0 THROUGH E7 -- SEE SECTION 4.2.5) TREATED AS CONTIGUOUS EXTENSIONS TO LOCAL STORE FOR THIS PURPOSE. SINCE THE MIX AND MOD BUSES MAY NOT BE CONNECTED TO AN EXTERNAL STORE REGISTER BEYOND THE PORT REGISTERS, THE FOLLOWING RULES APPLY:

1. WHEN FMIX CONTAINS A VALUE GREATER THAN 39 (CORRESPONDING TO E7, THE LAST PORT REGISTER), THE MIX BUS TAKES THE VALUE OF ALL ONES.
2. WHEN FMD CONTAINS A VALUE GREATER THAN 39, GATE M5 IS A NULL OPERATION.

4.3.2.3 SPECIAL F'S

THE NEXT SIX F-REGISTERS SERVE SPECIAL CONTROL FUNCTIONS, SOME OF WHICH HAVE BEEN MENTIONED PREVIOUSLY.

FACT: (AUXILLARY ACTION) FACT IS USED AS A SOURCE VALUE TO SPECIFY A VARIETY OF SPECIAL ACTION COMMANDS. THESE ARE DESCRIBED IN SECTION 5.8.2.

FLIV: (LOOP INPUT VALUE) FLIV IS AN F REGISTER AVAILABLE FOR USE IN "COUNTING" STEPS WITHIN A NANOPROGRAM. ANY F REGISTER CAN BE INCREMENTED OR DECREMENTED UNDER NANOPROGRAM CONTROL. IN ADDITION, A "SPECIAL CONDITION" CAN BE TESTED TO DETERMINE WHETHER OR NOT ANY SELECTED F REGISTER CONTAINS THE VALUE ZERO. FLIV IS AVAILABLE FOR PRIMARY USE IN THIS FASHION.

FMPC: THE CONTENTS OF FMPC, MODULO 4, SELECT ONE OF THE FOUR MICRO PROGRAM COUNTERS IN LOCAL STORE TO BE USED FOR MPC OPERATIONS (SEE SECTION 4.2.2). THE SELECTION IS ACCORDING TO:

FMPC (MOD 4)	MPC
0	R24
1	R25
2	R26
3	R27

FIDX: (INDEX) FIDX IS PARTITIONED INTO THREE FIELDS, EACH OF WHICH CONTROLS A DIFFERENT BASIC MACHINE FUNCTION. NUMBERING THE BITS 5, 4, 3, 2, 1, 0, THE FIELDS ARE AS FOLLOWS:

BIT 5: 16-BIT-MODE CONTROL; "1" FOR 16-BIT MODE, "0" FOR 18.

BITS 4 AND 3: MEMORY MODE CONTROL (SECTION 4.8);

"0" FOR NORMAL, "1" FOR READ-ONLY

BIT 4 - CONTROL STORE, BIT 3 - NANOSTORE

BITS 2, 1, 0: NANOSTORE PAGE INDEX, USED IN NANOSTORE

ADDRESSING UNDER MICRO CONTROL. (SECTIONS 4.5.3, 4.5.4)

FIST: (STATUS) THE SIX BITS OF FIST CONTAIN THE "GLOBAL CONDITIONS" MENTIONED IN SECTION 4.2.3. SINCE THE FIST BITS CAN BE TESTED BY NANOPRIMITIVE CONTROLS WHICH ARE INDEPENDENT OF THOSE USED

TO TEST THE "LOCAL CONDITIONS", AND SINCE ANY F-REGISTER CAN BE LOADED WITH A SIX-BIT QUANTITY, FIST MAY ALSO BE USED AS A GENERAL-PURPOSE BIT-TESTING FACILITY.

THE POSITIONS OF THE TEST BITS IN FIST ARE:

/	SHB	/	C	/	S	/	R	/	O	/	SLB	/

BIT	5		4		3		2		1		0	

FIPH: (PHANTOM) FIPH IS A SPECIAL F THAT GIVES THE NANOPROGRAMMER THE ABILITY TO TRANSFER A VALUE FROM A SOURCE AUX TO A DESTINATION AUX WITHOUT USING TWO T-STEPS AS WOULD BE REQUIRED WHEN GOING THROUGH AN F REGISTER, AND WITHOUT DESTROYING THE VALUE IN AN F REGISTER. THIS IS POSSIBLE BECAUSE FIPH IS NOT TRULY A REGISTER. HAVING NO DATA-STORING CAPABILITY, IT IS A DIRECT CONNECTION BETWEEN THE INPUT AND OUTPUT BUS STRUCTURES OF F-STORE. IF NANOPRIMITIVES ARE SIMULTANEOUSLY EXECUTED TO INPUT TO FIPH FROM A SIX-BIT SOURCE FIELD AND OUTPUT FROM FIPH INTO A SIX-BIT DESTINATION FIELD, THE RESULT IS A DIRECT TRANSFER FROM THE SOURCE FIELD TO THE DESTINATION FIELD. IF ONLY A LOAD INTO FIPH IS EXECUTED, THERE IS NO EFFECT; IF ONLY A GATE OUT OF FIPH IS EXECUTED, THE TRANSFERRED VALUE IS ZERO.

NOTE: TRANSFERS FROM SOURCE AUX TO DESTINATION AUX VIA FIPH MUST BE PLACED IN A STRETCHED T-STEP.

4.3.2.4 G'S

THE LAST TWELVE REGISTERS IN F-STORE ARE KNOWN AS G-FIELDS, OR G'S. THE G'S HAVE NO DIRECT DEDICATED MACHINE-CONTROL FUNCTIONS IN THEMSELVES, BUT ARE USED IN PROGRAMMING SYSTEMS TO STORE BACK-UP CONTROL INFORMATION, AS FOLLOWS:

FIRST, SINCE THE G'S ARE A PART OF F-STORE, ANY G MAY BE LOADED FROM ANY SOURCE AUX, OR READ INTO ANY DESTINATION AUX. (HENCE ONE USE OF G'S IS FOR TEMPORARY, OR SCRATCH, STORAGE IN SIX-BIT PROGRAMMING, WITHOUT INHIBITING THE USE OF MACHINE FUNCTIONS.)

SECOND, THE G'S HAVE THE SPECIAL PROPERTY THAT THEY ARE ALSO ADDRESSABLE AS SOURCE AUX FIELDS, AND HENCE MAY BE TRANSFERRED DIRECTLY TO ANY REGISTER IN F-STORE (INCLUDING G'S). THEREFORE THE G'S PROVIDE SPACE FOR THE PROGRAMMER TO STORE CONTROL INFORMATION THAT WILL SUBSEQUENTLY BE TRANSFERRED INTO (OR EXCHANGED WITH) BUS CONTROL F'S, SPECIAL F'S, AND/OR OTHER CONTROL REGISTERS (I.E., CERTAIN DESTINATION AUX FIELDS). IN THIS SENSE, THE G'S SERVE THE FUNCTION OF A "SECOND LEVEL" OF RESIDUAL CONTROL.

NOTE: AS A SOURCE, G'S MAY BE REFERRED TO AS G'S OR AS THEIR CORRESPONDING F'S, DEPENDING ON THE TRANSFER, E.G. G0 = F20. AS A DESTINATION THEY MAY ONLY BE F'S.

4.3.3 AUX FIELDS

AVAILABLE SOURCE AUX FIELDS (WHICH CAN BE LOADED INTO F-STORE) ARE:

C, A, B:	THE THREE SIX-BIT FIELDS OF R31, AS INTRODUCED IN SECTION 4.2.2.
KA, KB, KX, KT, KS:	SIX-BIT FIELDS FROM THE EXECUTING NANOWORD TO BE DISCUSSED BELOW AND IN SECTION 4.5.1. KA AND KB ORDINARILY ARE USED FOR CONSTANT AND/OR SCRATCH STORAGE.
G0 - G11:	THE 12 G'S (SEE SECTION 4.3.2.4).

THE FOLLOWING ADDITIONAL SOURCE AUX'S ARE NOT REGISTERS:

ALUF:	OUTPUT OF SIX-BIT ALU, TO BE DISCUSSED IN SECTION 4.3.4 (QM-1 OPTION).
ID ID:	A SIX-BIT IDENTIFICATION NUMBER ASSOCIATED WITH A DEVICE ON AN EXTERNAL PORT; SEE SECTIONS 4.6, 5.5.2.
INCF1, DECF1, INCF2, DECF2:	INCREMENTS AND DECREMENTS OF F-STORE ELEMENTS (SECTIONS 5.5.2, 5.6.6).

AVAILABLE DESTINATION AUX FIELDS (TO WHICH F-STORE CAN BE OUTPUT) ARE:

C, A, B:	SEE ABOVE.
KA, KB:	SEE ABOVE.
KALC:	A FIELD WHICH SPECIFIES THE OPERATION OF THE ALU (ALU CONTROL).
KSHC:	A FIELD WHICH SPECIFIES SHIFTING OPERATIONS (SHIFT CONTROL).
KSHA:	A FIELD WHICH SPECIFIES NUMBER OF POSITIONS TO SHIFT (SHIFT AMOUNT).
KS:	A SIX-BIT MASK FIELD ASSOCIATED WITH GLOBAL CONDITION (FIST) TESTING.
KT:	A SIX-BIT MASK FIELD ASSOCIATED WITH LOCAL CONDITION TESTING.
KX:	A SIX-BIT MASK FIELD ASSOCIATED WITH SPECIAL CONDITION TESTING.

KA, KB, KALC, KSHC, KSHA, KS, KT, AND KX ARE ALL SIX-BIT FIELDS IN THE EXECUTING NANOWORD (SEE SECTION 4.5.1).

4.3.4 ALUF (QM-1 OPTION) FURTHER SPECIFICATIONS TO BE ANNOUNCED)

A SIX-BIT ALU, SIMILAR IN CHARACTERISTICS TO THE 18-BIT ALU, OPERATES UNDER NANOPRIMITIVE CONTROL TO GENERATE ARITHMETIC AND LOGICAL FUNCTIONS FROM TWO SIX-BIT INPUTS.

THE LEFT INPUT TO "ALUF" IS ANY REGISTER IN F-STORE (SELECTED BY NANOPRIMITIVE CONTROLS); THE RIGHT INPUT IS SIMILARLY SELECTED FROM THE LIST:

GQ-G11, IO ID, B, KS, KT.

THE OUTPUT OF ALUF MAY BE LOADED INTO ANY REGISTER IN F-STORE.

IF THE ALUF IS NOT PRESENT, OPERATIONS DEFINED TO GATE ITS OUTPUT PRODUCE A ZERO VALUE.

4.4 TIMING

THE QM-1 CPU IS A SYNCHRONOUS DEVICE, DRIVEN BY A SINGLE MACHINE CLOCK WHOSE PERIOD IS 80 NANOSECONDS.

TO ALLOW THE HARDWARE-LEVEL USER INTIMATE ACCESS TO AND CONTROL OF QM-1 HARDWARE FACILITIES, NANOPROGRAM STEPS ARE EXECUTED AT THE MACHINE CLOCK RATE.

TO AVOID AMBIGUITY IN DISCUSSING THREE CLOSELY RELATED CONCEPTS, THE FOLLOWING TERMS ARE USED IN THIS MANUAL:

- T-PERIOD: A PERIOD OF ELAPSED TIME EQUAL TO THE CLOCK PERIOD; 80 NANOSECONDS.
- T-STEP: AN ELEMENTARY EVENT IN PROGRAM CONTROL; A SINGLE STEP OF NANOPROGRAM EXECUTION. A T-STEP CONSISTS OF THE SIMULTANEOUS (PARALLEL) EXECUTION OF SOME NUMBER OF NANOPRIMITIVE COMMANDS (NANOPRIMITIVES). A T-STEP GENERALLY OCCURS IN ONE T-PERIOD, BUT FOR CERTAIN PURPOSES IT MAY BE EXPANDED (BY THE "STRETCH" NANOPRIMITIVE) TO LAST FOR TWO T-PERIODS.
- T-VECTOR: A STRING OF BITS REPRESENTING A SET OF NANOPRIMITIVES TO BE EXECUTED CONCURRENTLY IN A SINGLE T-STEP. THE "ACTIVE T-VECTOR" CORRESPONDS TO THE "CURRENT T-STEP". (PROGRAM CONTROL IS PRESENTED IN DETAIL IN SECTION 4.5.)

WHEN THERE IS NO DANGER OF CONFUSING THE THREE CONCEPTS OF TIME (T-PERIOD), EVENT (T-STEP), AND PHYSICAL ENTITY (T-VECTOR), THE TERM "T" MAY BE USED; FOR EXAMPLE: "A 24-HOUR QM-1 WORKING DAY IS EQUIVALENT TO MORE THAN A TRILLION T'S." ("T" IS DERIVED FROM "TEE": TIME, EVENT, ENTITY.)

ALL NANOPRIMITIVES MAY BE CLASSIFIED AS EITHER "LEADING-EDGE" (LE) OR "TRAILING-EDGE" (TE), ACCORDING TO WHETHER THE FUNCTION THEY DEFINE TAKES EFFECT AT THE BEGINNING OR THE END, RESPECTIVELY, OF THE T-STEP IN WHICH THEY ARE EXECUTED. IN GENERAL, THE EFFECT OF TRAILING-EDGE NANOPRIMITIVES (THE LARGER CLASS) MAY BE CONSIDERED TO OCCUR AT THE END OF THE T-STEP IN

WHICH THEY OCCUR. FOR EXAMPLE, ALL NANOPRIMITIVES WHICH TRANSFER VALUES INTO REGISTERS (18-BIT OR SIX-BIT) ARE TRAILING-EDGE. LEADING-EDGE NANOPRIMITIVES, ON THE OTHER HAND, INITIATE PROCESSES WHICH HAVE A DURATION OF ONE OR MORE T-PERIODS; EXAMPLES ARE READ CS, MSGO. THE DURATION OF SUCH PROCESSES ARE MEASURED FROM THE BEGINNING OF THE T-STEP IN WHICH THEIR NANOPRIMITIVES ARE EXECUTED. THE "STRETCH" NANOPRIMITIVE SEPARATES THE LEADING EDGE OF A T-STEP FROM THE TRAILING EDGE BY ONE EXTRA T-PERIOD.

THE DIFFERENCE BETWEEN A T-STEP AND A T-PERIOD IS IMPORTANT WHEN BOTH LEADING-EDGE AND TRAILING-EDGE NANOPRIMITIVES ARE PROGRAMMED. FOR EXAMPLE, IF READ CS (LEADING-EDGE) AND GATE CS (TRAILING-EDGE) ARE PROGRAMMED IN THE SAME T-STEP, AND THAT T-STEP IS NOT STRETCHED, CONTROL STORE WILL NOT GENERATE THE NEW VALUE ON THE COD BUS IN TIME FOR THE GATE CS. IF, HOWEVER, THE T-STEP IS STRETCHED, THE VALUE GATED INTO LOCAL STORE WILL BE THAT GENERATED BY THE READ CS, SINCE THE TIME SPAN BETWEEN THE LEADING AND TRAILING EDGES OF A STRETCHED T-STEP IS TWO T-PERIODS, ENOUGH FOR A CONTROL STORE CYCLE. (TIMING CONSIDERATIONS FOR PROGRAMMING CONTROL STORE AND OTHER LEADING-EDGE OPERATIONS WILL BE DISCUSSED IN DETAIL IN SECTION 5.3)

ALL REGISTER TRANSFERS, BOTH 18-BIT AND SIX-BIT, ARE CONTROLLED BY TRAILING-EDGE NANOPRIMITIVES. SINCE THESE OPERATIONS ARE SYNCHRONOUS, THE SAME REGISTER EFFECTIVELY CAN BE LOADED AND READ IN THE SAME T-STEP ("SIMULTANEOUSLY"), WITHOUT LOSS OF DATA. GIVEN THE MODEL THAT REGISTER "REG" IS TO BE GATED TO REGISTER "DEST" AND LOADED FROM REGISTER "SOURCE" SIMULTANEOUSLY, THEN THE STATE BEFORE THE OPERATION IS THAT THE DATA FROM SOURCE IS PRESENT ON THE BUS FROM SOURCE TO REG, AND THE DATA IN REG IS PRESENT ON THE BUS FROM REG TO DEST; THE EFFECT OF THE NANOPRIMITIVES IS THEN TO LATCH ("CLOCK IN") THE BUS VALUES INTO REG AND DEST. PROPAGATION DELAYS ARE SUCH THAT DEST WILL HAVE LATCHED ITS NEW VALUE BEFORE THE NEW VALUE IN REG HAS TIME TO REACH THE BUS FROM REG TO DEST. IT IS QUITE IMPORTANT, HOWEVER, THAT THE NEW VALUES ARE ON THE BUSES AND READY FOR A REPEAT OPERATION WITHIN ONE T-PERIOD; THIS FACT IS FUNDAMENTAL IN NANOPROGRAMMING.

THE BASIC TIMING STRUCTURE OF THE QM-1 IS DERIVED NOT ONLY FROM HARDWARE CONSIDERATIONS, BUT ALSO FROM THE DESIGN OBJECTIVE OF BEING ABLE TO NANOPROGRAM CERTAIN OPERATIONS AND SEQUENCES OF OPERATIONS EFFICIENTLY. THE THREE MOST IMPORTANT SUCH SEQUENCES ARE LISTED BELOW (IN THE SYMBOLIC PROGRAM EXAMPLES, THE T-STEPS ARE NOT STRETCHED).

1. CLOSED LOOP THROUGH LOCAL STORE. FOR EXAMPLE, LET THE EOD AND EID BUSES BOTH CONNECT THE SAME LOCAL STORE REGISTER AND EXTERNAL STORE REGISTER (FEOD=FEID, FEOR=FEIR), AND THEN EXECUTE THE FOLLOWING T-STEPS:

TN: LOAD ES, GATE ES.
TN+1: LOAD ES, GATE ES.

THE RESULT IS A DOUBLE EXCHANGE (FINAL STATUS = INITIAL STATUS) OF THE CONTENTS OF THE TWO REGISTERS.

2. CLOSED LOOP THROUGH F-STORE; E. G. :

TN: AUX(X)---->F(Y), F(Y)---->AUX(X).
TN+1: AUX(X)---->F(Y), F(Y)---->AUX(X).

THE RESULT IS A DOUBLE EXCHANGE (FINAL STATUS = INITIAL STATUS) OF THE CONTENTS OF AUX(X) AND F(Y).

3. BUS SETTING IMMEDIATELY PRIOR TO BUS USE; E. G. :

TN: AUX(X)---->FSOD.
TN+1: GATE SH.

THE SHIFTER OUTPUT IS GATED INTO THAT LOCAL STORE REGISTER "POINTED TO" BY THE NUMBER CONTAINED IN AUX(X) AT THE BEGINNING OF T-STEP T(N).

THE ABOVE THREE TYPES OF SEQUENCES ARE MUTUALLY CONSISTENT. FOR EXAMPLE, THE THIRD ILLUSTRATION COULD BE EXPANDED TO INCLUDE THE T-STEP:

TN-1: F(Y)---->AUX(X).

WHICH WOULD SET AUX(X) IN TIME FOR THE DESCRIBED SEQUENCE TO OCCUR WITH THE LOCAL STORE REGISTER NUMBER SPECIFIED BY F(Y).

TO ACHIEVE THESE OBJECTIVES, THE MACHINE CLOCK SIGNAL THAT CONTROLS SIX-BIT OPERATIONS (THE F CLOCK) IS PHASED DIFFERENTLY FROM THE SIGNAL THAT CONTROLS 18-BIT OPERATIONS (THE R CLOCK). BOTH ARE DERIVED FROM THE SAME SINGLE MACHINE CLOCK THAT CONTROLS THE T-VECTOR (THE T CLOCK). THIS PHASING IS NORMALLY TRANSPARENT TO THE NANOPROGRAMMER, AND MUST BE CONSIDERED MAINLY WHEN DEALING WITH THE INTERFACE BETWEEN THE SIX-BIT ARCHITECTURE AND THE 18-BIT ARCHITECTURE WHICH EXISTS IN R31; THE NECESSARY PROGRAMMING RULES ARE SPECIFIED IN DETAIL IN SECTION 5.3.5.

4.5.1 CONTROL MATRIX

AT ANY GIVEN TIME, THE CONTROL MATRIX CONTAINS 360 BITS, CORRESPONDING TO THE 360 BITS IN A NANOWORD. THIS INCLUDES THE FOUR T-VECTORS, ONE OF WHICH IS ACTIVE (72 BITS IN EACH T-VECTOR) -- AND A 72-BIT ENTITY, THE "K-VECTOR". (THE AUX REGISTERS KA, KB, KALC, KSHC, KSHA, KS, KT, AND KX ARE ALL PORTIONS OF THE K-VECTOR.) SINCE THE CONTROL MATRIX BITS CORRESPOND TO SOME 360-BIT WORD ("NANOWORD") IN NANOSTORE, THEY MAY BE REFERRED TO AS THE "ACTIVE

NANOWORD". BECAUSE OF THIS CORRESPONDENCE, EVERY WORD IN NANOSTORE IS LOGICALLY PARTITIONED AS FOLLOWS:

[K-VECTOR] [T-VECTOR(1)] [T-VECTOR(2)] [T-VECTOR(3)] [T-VECTOR(4)]

(OR, BRIEFLY: K, T1, T2, T3, T4 -- HIGH-ORDER TO LOW-ORDER)

MECHANISMS ARE PROVIDED FOR SELECTING A NANOWORD, FETCHING THAT WORD FROM NANOSTORE, AND LOADING IT INTO THE CONTROL MATRIX. WHEN THE NANOWORD IS LOADED INTO THE CONTROL MATRIX, ITS FIRST T-VECTOR (T1) IMMEDIATELY BECOMES THE ACTIVE T-VECTOR (AND ITS K-VECTOR BECOMES ACTIVE).

THE NORMAL OPERATION OF THE CONTROL MATRIX ACTIVATES THE FOUR T-VECTORS IN SUCCESSION AND CIRCULARLY, WITH NO LOSS OF TIME BETWEEN ACTIVATIONS: T1, T2, T3, T4, T1, ETC. UNLESS A SPECIAL HIGH-PRIORITY FACILITY (PROGRAM CHECK) INTERRUPTS, THIS SEQUENCE CONTINUES UNTIL CERTAIN PROGRAM-CONTROL NANOPRIMITIVES ARE EXECUTED. THESE NANOPRIMITIVES CAN BE PROGRAMMED TO EXECUTE CONDITIONALLY, SO THAT THE NANOPROGRAMMER MAY CREATE A USEFUL LOOP IN A SINGLE NANOWORD. (FOR EXAMPLE, THE F ZERO TEST MAY BE USED; SEE SECTION 5.7.) IF THE PROGRAMMER DOES NOT NEED SUCH LOOPING, THEN THE SEQUENCE MAY OF COURSE BE BROKEN AFTER THE FIRST ACTIVATION OF T4 (OR EARLIER, IF DESIRED).

FOR PROTECTION AGAINST INFINITE LOOPING, A CONTROL MATRIX TIME-OUT FACILITY BREAKS THE LOOP AND GENERATES A PROGRAM CHECK IF THE SAME NANOWORD CIRCULATES IN THE CONTROL MATRIX FOR APPROXIMATELY ONE SECOND (MORE THAN 12 MILLION T-PERIODS).

THE TWO PROGRAM CONTROL NANOPRIMITIVES OF IMMEDIATE INTEREST ARE "SKIP" AND "GATE NS". EITHER CAN BE EXECUTED CONDITIONALLY, ACCORDING TO THE T-VECTOR TEST FACILITIES SPECIFIED IN SECTION 5.7. THE BIT STRUCTURE IN THE T-VECTOR IS SUCH THAT SKIP AND GATE NS ARE MUTUALLY EXCLUSIVE IN THE SAME T-STEP. HOWEVER, ANOTHER NANOPRIMITIVE "GATE NS UNCONDITIONALLY" IS PROVIDED TO AVOID THIS RESTRICTION AND TO PERMIT CONDITIONALLY SKIPPING T1 OF THE NEXT NANOWORD.

SKIP, WHEN EXECUTED, MODIFIES CONTROL MATRIX OPERATION SO THAT THE NEXT T-STEP IN SUCCESSION IS SKIPPED OVER; ACTIVATION OF THE SUCCEEDING T-VECTOR IS INHIBITED. THE SKIPPED T-STEP CONSUMES ONE T-PERIOD OF TIME (WHETHER STRETCHED OR NOT), WHICH SHOULD BE OBSERVED WHEN LEADING-EDGE PROCESSES ARE PROGRAMMED; THE EFFECT IS THE SAME AS IF THE SUCCEEDING T-VECTOR WERE ACTIVATED, BUT WITH ALL ITS SPECIFIED NANOPRIMITIVES (INCLUDING STRETCH) "TURNED OFF". NOTE THAT A SKIP EXECUTED IN T3 RESULTS IN T1 BEING THE NEXT

T-VECTOR ACTIVATED; A SKIP IN T4 GOES TO T2.

GATE NS IS A TRAILING-EDGE NANOPRIMITIVE WHICH, WHEN EXECUTED, CAUSES THE CONTROL MATRIX TO BE LOADED WITH THE NANOWORD RESULTING FROM THE MOST RECENTLY COMPLETED NANOSTORE ACCESS. THE SUCCESSOR TO THE T-STEP IN WHICH GATE NS IS EXECUTED IS GENERATED BY THE FIRST T-VECTOR (T1) OF THE NEWLY GATED NANOWORD; NO TIME IS LOST IN THE TRANSITION. THE K-VECTOR IS ALSO LOADED FROM THE K-VECTOR OF THE NANOWORD AS A RESULT OF THE GATE NS.

NOTE: ANY SIX-BIT TRANSFERS INTO K-VECTOR AUX FIELDS COMMANDED IN A T-STEP IN WHICH GATE NS IS ALSO EXECUTED RESULT IN UNDEFINED VALUES, UNLESS THE "HOLD" CONTROL IS ON (SEE SECTION 4.5.2.3).

TO SUPPLY A NANOWORD FOR GATING INTO THE CONTROL MATRIX, THE LEADING-EDGE NANOPRIMITIVE READ NS IS EXECUTED. NANOSTORE COMPLETES THE READ OPERATION WITHIN TWO T-PERIODS (BUT NOT WITHIN ONE T-PERIOD); HENCE EITHER OF THE FOLLOWING PROGRAM EXAMPLES ILLUSTRATE A SUCCESSFUL SHORTEST-TIME SEQUENCE OF THE OPERATIONS READ NS, GATE NS:

- A) TN: READ NS. (NOT STRETCHED)
 TN+1: GATE NS. (NOT STRETCHED)
- B) TN: STRETCH, READ NS, GATE NS.

THE READ NS NANOPRIMITIVE HAS A SECONDARY EFFECT, INVOLVING THE "NANOBRANCH" FACILITY; THIS IS DISCUSSED IN SECTION 4.5.2.3.

4.5.2 NANOSTORE ADDRESSING

4.5.2.1 PRIORITY SELECT

WHEN READ NS IS EXECUTED, A PRIORITY-SELECT MECHANISM SUPPLIES THE ACTUAL NANOSTORE ADDRESS FROM A LIST OF POTENTIAL ADDRESSES. EACH POTENTIAL ADDRESS VALUE IS 10 BITS WIDE, SINCE NANOSTORE MAY CONTAIN AS MANY AS 1024 NANOWORDS.

NANOSTORE IS AVAILABLE IN 256-WORD BLOCKS, AND CAN BE ARRANGED SO THAT ANY OF EIGHT POSSIBLE 128-WORD "PAGES" IS FULL, HALF-FULL, OR EMPTY. EXECUTION OF READ NS FROM A NONEXISTANT LOCATION GENERATES A ZERO NANOWORD; IF SUCH A WORD IS LOADED INTO THE CONTROL MATRIX, NO NANOPRIMITIVE OPERATIONS ARE INVOKED, AND A CONTROL MATRIX TIME-OUT PROGRAM CHECK EVENTUALLY RESULTS.

EACH SOURCE OF POTENTIAL NANOSTORE ADDRESS HAS A FIXED PRIORITY RELATIVE TO THE OTHER SOURCES; FURTHERMORE, AN ACTIVE/INACTIVE STATUS IS ASSOCIATED WITH EACH SOURCE AT ANY GIVEN TIME. WHEN INVOKED, THE PRIORITY-SELECT MECHANISM SELECTS THE ADDRESS FROM THE HIGHEST-PRIORITY SOURCE THAT IS CURRENTLY ACTIVE AND SUPPLIES IT TO NANOSTORE. IF THE CORRESPONDING NANOWORD IS THEN EXECUTED (GATE NS OCCURS BEFORE THE NEXT READ NS), THE ADDRESS SOURCE IS RESET TO INACTIVE STATUS. THE SOURCE WITH LOWEST FIXED PRIORITY IS THE NANOPROGRAM COUNTER (NPC); THIS IS DEFINED AS PERMANENTLY ACTIVE, AND CAN BE CONSIDERED A DEFAULT.

THE PROCESS CAN BE MODELED BY A LIST OF ACTIVATION FLAGS, EACH ASSOCIATED WITH A NANOSTORE-ADDRESS SOURCE, ORDERED BY THE PRIORITY OF THE SOURCES; OPERATION OF THE PRIORITY-SELECT MECHANISM IS EQUIVALENT TO READING DOWN THIS LIST, FROM HIGH-PRIORITY TO LOW-PRIORITY, UNTIL THE FIRST ACTIVE FLAG IS ENCOUNTERED. THE ADDRESS ASSOCIATED WITH THAT FLAG IS THEN SUPPLIED TO NANOSTORE, AND THE FLAG IS TURNED OFF (INACTIVE) UPON SUCCESSFUL USE (GATE NS) OF THE NANOWORD FETCHED FROM THAT LOCATION.

FIGURE 4.5.2.1A ILLUSTRATES THE MODEL AND IDENTIFIES THE VARIOUS ADDRESS SOURCES, TO BE DISCUSSED IN THE FOLLOWING SECTIONS.

FIGURE 4.5.2.1A

PRIORITY SELECTION OF NANOSTORE ADDRESS

FLAGS	SOURCE (HIGHEST PRIORITY AT TOP)
-------	----------------------------------

[]	PROGRAM CHECK
[]	NANOBANCH
	--
[]	I
.	I
.	I INTERRUPTS (MAXIMUM 30)
.	I
[]	I
	--
[X]	NANOPROGRAM COUNTER (NPC) (PERMANENTLY ACTIVE)

UPON EXECUTION OF READ NS, THE MECHANISM READS DOWN FROM THE TOP TO FIRST FLAG THAT IS ACTIVE [X]. ASSOCIATED 10 BIT ADDRESS IS SENT TO NANOSTORE. IF THE CORRESPONDING NANOWORD IS EXECUTED (GATE NS), FLAG IS TURNED OFF.

NOTE: THE FLAG ASSOCIATED WITH NANOBANCH IS TREATED IN A SPECIAL MANNER TO BE DESCRIBED IN SECTION 4.5.2.3.

4.5.2.2 PROGRAM CHECK

WHEN A PROGRAM CHECK OCCURS, THE FOLLOWING IS AUTOMATICALLY DONE:

1. EXECUTION OF THE ACTIVE NANOWORD IS TERMINATED.
2. THE APPROPRIATE BIT IS SET IN THE PROGRAM CHECK STATUS FIELDS (HIGH 3 BITS OF E20, HIGH 3 BITS OF E21).
3. THE MEMORY MODE BITS IN FIDX ARE FORCED TO "1", ALLOWING ACCESS TO READ-ONLY NANOSTORE (RONS) AND READ-ONLY CONTROL STORE (ROCS) (THESE WILL BE DISCUSSED IN SECTION 4.8).
4. THE CONTENTS OF RONS[0] ARE LOADED INTO THE CONTROL MATRIX TO BEGIN EXECUTION OF THE SERVICE PROGRAM.

SINCE THE ENTRY POINT OF RONS[0] IS SHARED BY THE PROGRAM CHECK SERVICE PROGRAM AND THE MACHINE START PROGRAM (SEE SECTION 4.8), THE COMMON PROGRAM MUST TEST FOR PROGRAM CHECK STATUS FIELDS OF ZEROS (CLEARED BY THE MACHINE START PUSHBUTTON) TO DETERMINE THE NATURE OF ITS INVOCATION. A "SPECIAL CONDITION", SET TO "TRUE" IF ANY OF THE PROGRAM CHECK STATUS BITS ARE ON, CAN BE TESTED FOR THIS PURPOSE (SEE SECTION 5.7.1 AND 5.7.2).

PROGRAM CHECK STATUS BITS ARE AS FOLLOWS:

E20, BIT 17:	MS PARITY ERROR
E20, BIT 16:	MS ADDRESS VIOLATION
E20, BIT 15:	(RESERVED FOR FUTURE EXPANSION)
E21, BIT 17:	ILLEGAL MICRO-OPCODE ENTRY
E21, BIT 16:	(RESERVED FOR FUTURE EXPANSION)
E21, BIT 15:	TIME-OUT (CONTROL MATRIX OR MICROINSTRUCTION)

4. 5. 2. 3 NANOBANCH

THE NANOBANCH FACILITY IS ONE MEANS OF CONTINUING A NANOPROGRAM BEYOND ONE NANOWORD. DUE TO THE HIGH PRIORITY GIVEN TO THE NANOBANCH OPERATION, A BRANCH-CONNECTED NANOPROGRAM IS NEVER INTERRUPTIBLE EXCEPT BY PROGRAM CHECK.

THE SOURCE OF THE NANOBANCH ADDRESS IS A 10-BIT FIELD, KN, WITHIN THE ACTIVE K-VECTOR.

CONTROL OF THE NANOBANCH ACTIVITY STATUS FOR PRIORITY SELECTION IS ACCOMPLISHED THROUGH THE BRANCH BIT IN THE ACTIVE K-VECTOR, IN CONJUNCTION WITH THE READ NS NANOPRIMITIVE:

EACH TIME A READ NS IS EXECUTED, THE BRANCH BIT IS TESTED. IF ACTIVE, THE NANOBANCH ADDRESS IS TAKEN. IF INACTIVE, THE NANOSTORE ADDRESS IS TAKEN FROM ONE OF THE LOWER PRIORITY SOURCES AS DESCRIBED IN SECTION 4. 5. 2. 1. THUS, THE BRANCH BIT SERVES AS THE ACTIVATION FLAG FOR THE SELECTION OF THE NANOBANCH ADDRESS.

THE INITIAL CONDITION OF THE BRANCH BIT IS DETERMINED BY ITS SETTING IN THE NANOWORD GATED INTO THE CONTROL MATRIX. IF SET, BRANCH IS ACTIVE AS SOON AS THE NANOWORD (I.E., THE ONE CONTAINING THE BRANCH BIT) IS LOADED INTO THE CONTROL MATRIX. THE STATE OF ANOTHER BIT, THE "ALTERNATE" BIT IN THE ACTIVE K VECTOR DETERMINES THE FUTURE CONDITION OF THE BRANCH FLAG. WHEN ALTERNATE IS NOT SET, THE BRANCH BIT RETAINS ITS INITIAL STATUS. HOWEVER, WHEN THE ALTERNATE BIT IS SET, EVERY EXECUTION OF READ NS IN THE ACTIVE NANOWORD ACQUIRES THE SECONDARY FUNCTION OF INVERTING (COMPLEMENTING) THE BRANCH ACTIVITY FLAG AFTER INITIATING THE NANOSTORE READ OPERATION.

THUS, THE NANOPROGRAMMER CAN SPECIFY FOUR POSSIBLE SETTINGS OF THESE TWO BITS TO CONTROL THE SELECTION OF THE NANOBANCH ADDRESS:

ALTERNATE	BRANCH	ACTION
0	0	NANOBANCH ADDRESS NEVER USED BY READ NS
0	1	NANOBANCH ADDRESS ALWAYS USED BY READ NS
1	0	2ND, 4TH, 6TH, ETC. READ NS USES NANOBANCH ADDRESS
1	1	1ST, 3RD, 5TH, ETC. READ NS USES NANOBANCH ADDRESS

NOTE THAT THIS CONTROL IS ALWAYS DETERMINED ON THE READ NS EXECUTION AND NOT ON THE GATE NS AS FOR OTHER ACTIVATION FLAGS. THUS THE SKIP AND/OR GATE NS FACILITIES MAY BE USED TO EFFECT A VARIETY OF CONDITIONAL NANOBANCHING.

4.5.2.4 EXTERNAL INTERRUPTS

INTERRUPTS ARE SIGNALS WHICH CAN NOTIFY THE PROGRAM OF THE OCCURRENCE OF EVENTS EXTERNAL TO THE QM-1. A MAXIMUM OF 30 SUCH SIGNALS ARE DIRECTLY DETECTABLE BY QM-1 HARDWARE.

THE 30 INTERRUPTS ARE ORDERED BY PRIORITY LEVEL FOR NANOSTORE ADDRESS SELECTION AND ARE LABELED LEVEL 2 THROUGH LEVEL 31; LEVEL 2 IS HIGHEST-PRIORITY (IMMEDIATELY BELOW NANOBANCH), AND LEVEL 31 IS LOWEST PRIORITY (IMMEDIATELY ABOVE THE NANOPROGRAM COUNTER).

ASSIGNMENT OF LEVELS TO SIGNAL LINES IS PERFORMED BY NANODATA AT INSTALLATION TIME ACCORDING TO USER SPECIFICATIONS. A TYPICAL ASSIGNMENT IS SHOWN IN THE ADJOINING FIGURE.

FOR AN INTERRUPT TO BECOME ACTIVE FOR PRIORITY SELECTION, IT MUST BE "LATCHED", "ENABLED", "PENDING" AND "ALLOWED".

AN INTERRUPT LEVEL IS LATCHED WHEN A 50 NS. PULSE IS SENSED ON ITS SIGNAL LINE.

INTERRUPT LEVELS ARE INDIVIDUALLY ENABLED BY THE "1" STATE OF THE CORRESPONDING INTERRUPT ENABLE BIT. THESE 30 BITS ARE STORED IN EXTERNAL STORE REGISTERS 18 AND 19 (SEE SECTION 4.2.5).

EVERY ENABLED LEVEL IS TESTED FOR THE PRESENCE OF A LATCHED INTERRUPT SIGNAL BY EACH EXECUTION OF GATE NS. IF THIS TEST SUCCEEDS, THE LEVEL IS SET TO PENDING STATUS. A PENDING FLAG FOR THE PARTICULAR INTERRUPT LEVEL IS SET IN THE CORRESPONDING INTERRUPT PENDING BIT OF EXTERNAL STORE REGISTERS 20 AND 21 (SEE SECTION 4.2.5). ONCE A LEVEL IS PENDING, IT REMAINS IN THAT STATE UNTIL THE PRIORITY-SELECT MECHANISM EVENTUALLY SELECTS THE NANOSTORE ADDRESS CORRESPONDING TO THAT INTERRUPT LEVEL, AND THE ASSOCIATED NANOWORD IS LOADED INTO THE CONTROL MATRIX TO BEGIN THE SERVICE PROGRAM; AT THAT TIME THE LEVEL IS ALSO UNLATCHED.

A PENDING INTERRUPT LEVEL AUTOMATICALLY BECOMES ACTIVE FOR PRIORITY SELECTION IF AND ONLY IF ITS ASSOCIATED ALLOW INTERRUPT BIT IN THE ACTIVE

INTERRUPT LEVEL ASSIGNMENTS

CHANNEL NUMBER	LEVEL ASSIGNED		
	DATA IN	DATA OUT	STATUS
0	2	3	22
1	4	5	23
2	6	7	24
3	8	9	25
4	10	12	26
5	13	11	27
6	14	15	28
7	16	17	29

LEVELS 18-21, 30 AND 31 MAY BE ASSIGNED TO OTHER EXTERNAL SIGNALS.

LEVELS 2-11 - NANO INTERRUPTS.

LEVELS 12-31 - MICRO INTERRUPTS.

K-VECTOR IS "1" WHEN THE PRIORITY-SELECT MECHANISM IS INVOKED BY READ NS. THERE ARE TWO SUCH BITS ("ALLOW NANO INTERRUPT" AND "ALLOW MICRO INTERRUPT"). THE HIGH 10 INTERRUPT LEVELS (2-12) ARE DESIGNATED AS NANO INTERRUPT LEVELS. THE REMAINING 20 LEVELS (12-31) ARE DESIGNATED AS MICRO INTERRUPT LEVELS.

IF A NANOBANCH IS NOT TAKEN AT THE END OF EXECUTING A NANOWORD, AND IF NO INTERRUPTS ARE ACTIVE, THE PRIORITY-SELECT MECHANISM GIVES CONTROL TO THE NANOPROGRAM COUNTER (SEE FIGURE 4.5.2.1A); HENCE THE ALLOW INTERRUPT BITS ARE A FACILITY THE PROGRAMMER CAN USE TO INSURE THAT A CHAIN OF NANOWORD EXECUTIONS INVOKED THROUGH THE NPC IS NOT INTERRUPTED. THIS SUBJECT IS FURTHER DISCUSSED IN SECTION 4.5.4.2.

ALL I/O INTERRUPTS MAY BE BLOCKED FROM PRIORITY SELECTION BY DISABLING I/O INTERRUPTS WITH THE AUXILLIARY ACTION "DISABLE" COMMAND AS DESCRIBED IN SECTION 5.8.2. THE ALLOW INTERRUPT BITS ARE THEN IGNORED AND NO I/O INTERRUPTS WILL BE ACCEPTED. FOLLOWING THE AUXILLARY ACTION "ENABLE" COMMAND, ALL BLOCKED I/O INTERRUPTS AGAIN BECOME ELIGIBLE FOR PRIORITY SELECTION, ASSUMING ALL OTHER PREREQUISITES EXIST, AS DESCRIBED ABOVE.

THE NANOSTORE ADDRESSES ASSOCIATED WITH THE INTERRUPT LEVELS ARE GENERATED FROM SIX-BIT FIELDS IN TEN EXTERNAL STORE REGISTERS 22 - 31 (SEE SECTION 4.2.5). THE MAPPING OF THE SIX-BIT FIELD INTO THE 10-BIT ADDRESS IS AS

MAPPING OF INTERRUPT ADDRESSES

SIX-BIT FIELD: "ABCDEF"
TEN-BIT ADDRESS: "0AB000CDEF"

SHOWN ON THE RIGHT. THIS MAPPING PERMITS UP TO 16 INTERRUPT ENTRIES IN EACH PAGE OF NANOSTORE. TO CONSERVE ENTRY POINTS IN NANOSTORE, SEVERAL INTERRUPT LEVELS MAY BE ASSIGNED THE SAME ADDRESS IN NANOSTORE BY PLACING THE COMMON ADDRESS IN THE APPROPRIATE POSITIONS IN EXTERNAL STORE REGISTERS 22 - 31.

FINALLY, A FACILITY EXISTS FOR THE PROGRAMMER TO "GENERATE" INTERRUPTS (SIMULATE EXTERNAL INTERRUPTS) AND ALSO TO "CLEAR" INTERRUPT LATCHES. ONE INTERRUPT LEVEL CAN BE SO AFFECTED IN A NANOWORD. WHEN THE "GENERATE INT" BIT IS SET IN THE ACTIVE K-VECTOR, THE SELECTED LEVEL IS LATCHED OR UNLATCHED AT THE BEGINNING OF T-PERIOD 1: IF THE "GENERATE" OPTION IS USED, THE LEVEL BECOMES LATCHED, PENDING, AND ACTIVE IN TIME TO BE SELECTED FOR EXECUTION BY A READ NS EXECUTED IN T-PERIOD 3 OR LATER; IF THE "CLEAR" OPTION IS USED, THE LEVEL IS UNLATCHED IMMEDIATELY UPON ACTIVATION OF THE CURRENT NANOWORD. THE MECHANISM FOR SELECTING THE LEVEL AND THE "GENERATE" VS. "CLEAR" OPTION IS PRESENTED IN SECTION 5.8.1.

OTHER EXTERNAL INTERFACE FACILITIES ARE DISCUSSED IN SECTION 4.6.

4.5.3 NANOPROGRAM COUNTER

4.5.3.1 GENERAL

WHEN AN NANOBANCH IS NOT TAKEN AND NO INTERRUPTS ARE ACTIVE, THE PRIORITY-SELECT MECHANISM SUPPLIES AN ADDRESS TO NANOSTORE FROM THE NANOPROGRAM COUNTER, THE LOWEST ELEMENT ON THE PRIORITY LIST.

THE NPC IS A 10-BIT REGISTER WHICH CHANGES VALUE ONLY AS A RESULT OF NANOPRIMITIVE COMMANDS. THE FOLLOWING (MUTUALLY EXCLUSIVE) NPC CONTROL OPERATIONS ARE AVAILABLE IN THE T-VECTOR (ALL TRAILING-EDGE):

- LOAD NPC (CS)
- LOAD NPC (KN)
- LOAD NPC (SEQUENCE)

THE FIRST OPERATION -- LOAD NPC (CS) -- INVOLVES MICROINSTRUCTION EXECUTION, SINCE THE ADDRESS IS A CONTROL STORE OPCODE. THIS IS DISCUSSED IN THE NEXT SECTION.

THE LOAD NPC (KN) OPERATION LOADS THE NANOPROGRAM COUNTER FROM THE KN FIELD IN THE ACTIVE K-VECTOR. THUS AN EXECUTING NANOWORD CAN TRANSFER NANOPROGRAM EXECUTION TO NS(KN) EITHER DIRECTLY (NANOBANCH) OR THROUGH NPC (NPC BRANCH). THE INTERRELATIONSHIP OF THESE TWO FACILITIES IS DISCUSSED IN SECTION 4.5.4.

THE LOAD NPC (SEQUENCE) OPERATION ADDS ONE (MODULO 1024) TO THE CONTENTS OF THE NPC. THUS A NANOPROGRAM EXECUTING AT AN NPC-SPECIFIED NANOSTORE ADDRESS CAN CONVENIENTLY CONTINUE EXECUTION THROUGH SEQUENTIAL NANOSTORE LOCATIONS (NANOSEQUENCING).

4.5.3.2 MICROINSTRUCTION EXECUTION

ONE OF THE MOST IMPORTANT MODES OF PROGRAM CONTROL IS THE INVOCATION OF A NANOPROGRAM BY A MICROINSTRUCTION; THE OPERATION CODE OF A MACHINE MICRO-INSTRUCTION, EXTRACTED FROM CONTROL STORE, IS USED TO SELECT THE NANOSTORE ENTRY ADDRESS OF THE NANOPROGRAM (OF ONE OR MORE NANOWORDS) WHOSE EXECUTION DEFINES THAT MICROINSTRUCTION.

WHEN THE NANOPRIMITIVE COMMAND LOAD NPC (CS) IS EXECUTED, THE FOLLOWING OCCURS (TRAILING-EDGE):

- A) THE HIGH-ORDER THREE BITS OF NPC ARE LOADED WITH THE NANOSTORE PAGE INDEX FROM THE LOW-ORDER THREE BITS OF FIDX, A SPECIAL F-REGISTER (SEE SECTION 4.3.2.3);
- B) THE LOW-ORDER SEVEN BITS OF NPC ARE LOADED FROM THE HIGH-ORDER SEVEN BITS OF THE COD BUS;
- C) THE LOW-ORDER ELEVEN BITS OF THE COD BUS ARE SAVED IN A DEDICATED REGISTER.

THE NANOPRIMITIVE "LOAD R31" IS AVAILABLE TO CAUSE THE FOLLOWING ACTION:

- A) THE HIGH-ORDER SEVEN BITS OF R31 ARE CLEARED TO ZEROS; AND
- B) THE LOW-ORDER ELEVEN BITS OF R31 ARE LOADED WITH THE SAVED ELEVEN LOW-ORDER BITS OF COD (THIS IS THE PARAMETER PART OF THE MACHINE MICROINSTRUCTION).

THIS EVENT IS CONCURRENT WITH SIX-BIT TRANSFERS EXECUTED IN THE LAST T-STEP OF THE PREVIOUS ACTIVE NANOWORD; THUS THE NEW CONTENTS OF THE C, A, AND B FIELDS IN R31 ARE AVAILABLE FOR GATING TO F-STORE IN THE FIRST T-STEP OF THE MICROINSTRUCTION, IF DESIRED (SEE SECTION 5.3.5).

THE HIGH-ORDER SEVEN BITS OF A MACHINE MICROINSTRUCTION ARE THUS DEFINED AS THE MICRO-OPCODE, AND PROVIDE THE MICROPROGRAMMER WITH A MAXIMUM OF 128 MICROINSTRUCTIONS SUPPORTED BY A PAGE OF NANOSTORE; I. E., 128 NPC-ADDRESSES ARE POSSIBLE UNDER A GIVEN VALUE OF THE NANOSTORE PAGE INDEX IN FIDX. DIFFERENT NANOSTORE PAGES MAY BE USED TO DEFINE DIFFERENT MICRO-MACHINES, EXTEND THE MICROINSTRUCTION SET OF A GIVEN MICRO-MACHINE, IMPLEMENT DIFFERENT MACHINE STATES, AND/OR CONTAIN CONTINUATIONS OF NANOPROGRAMS FROM ANOTHER PAGE. FOR CONVENIENCE IN MICROINSTRUCTION SEQUENCING, ONE OF THE LOCAL STORE MPC'S WILL NORMALLY BE USED IN ADDRESSING CONTROL STORE TO GENERATE THE MACHINE MICROINSTRUCTION ON THE COD BUS (SEE SECTION 4.2.4).

4.5.4.2 INTERRUPTIBILITY

ONE SUGGESTED MODE OF INTERRUPTIBILITY IS TO ALLOW LOW-PRIORITY INTERRUPT LEVELS TO TAKE CONTROL ONLY BETWEEN MICROINSTRUCTIONS; THIS PLAN IS ENFORCED BY SETTING THE ALLOW MICRO INTERRUPT BIT (IN THE K-VECTOR) ONLY IN THE EXITING NANOWORD OF A NANOPROGRAM (I. E., THAT WORD WHICH IS ORDINARILY SUCCEEDED BY A NANOPROGRAM INVOKED BY THE NEXT MACHINE MICROINSTRUCTION). IN THIS MODE, THE INTERRUPT-SERVICE NANOPROGRAMS ARE FREE TO MAKE USE OF THE NPC TO ESTABLISH A MAINLINE, AND THUS CALL SUBROUTINES, TRANSFER TO MICROINSTRUCTION CONTROL, ETC.; THE LOWER-PRIORITY INTERRUPTS ARE MORE LIKELY TO REQUIRE SUCH SERVICE (E. G., END OF IO OPERATION).

A SUGGESTED PARALLEL MODE OF INTERRUPTIBILITY IS TO ALLOW MAINLINE NANOPROGRAMS TO BE INTERRUPTED BETWEEN (SOME) NANOWORDS, USING THE ALLOW NANO INTERRUPT BIT. (NOTE: BRANCH-CONNECTED NANOPROGRAMS, INCLUDING SUB-NANOPROGRAMS AS DEFINED IN SECTION 4.5.4.1, CANNOT BE INTERRUPTED BECAUSE OF THE HIGH FETCH-PRIORITY OF NANOBRANCH.) WHEN ALLOWING THE MAINLINE NANOPROGRAM TO BE INTERRUPTED, THE INTERRUPT-SERVICE NANOWORDS MUST PROCEED BY NANOBRANCH ONLY, SINCE ANY OTHER TECHNIQUE WOULD DESTROY THE VALUE OF THE NPC AND HENCE BREAK THE LINK FOR RETURNING TO THE INTERRUPTED (MAINLINE) NANOPROGRAM. THE HIGHER-PRIORITY INTERRUPTS ARE MORE LIKELY TO BE SERVICEABLE BY THIS KIND OF PROGRAM (E. G., SINGLE-WORD TRANSFER IN A DATA STREAM).

ANY PROGRAMMING STRUCTURE MUST ALLOCATE VARIOUS MACHINE RESOURCES TO THE VARIOUS LEVELS OF PROGRAM CONTROL. FOR EXAMPLE, PROGRAMMING CONVENTIONS COULD BE ESTABLISHED SUCH THAT BUS CONTROLS ARE UNDEFINED BETWEEN MICROINSTRUCTIONS (FOR FREE USABILITY BY INTERRUPT SERVICE ROUTINES), BUT ARE EXPECTED TO HOLD BETWEEN MAINLINE NANOWORDS; THAT CERTAIN G'S DO NOT CHANGE VALUE BETWEEN MICROINSTRUCTIONS; ETC.

4.5.4.3 HOLD

FOR CONVENIENCE IN PROGRAMMING ACROSS NANOWORDS, THE "HOLD" AND "HOLD 2" BITS IN THE ACTIVE K-VECTOR ARE PROVIDED. THEY ALLOW THE NANOPROGRAMMER TO RETAIN VARIOUS CONTROL VALUES IN THE K VECTOR PORTION OF THE CONTROL MATRIX DURING THE TRANSITION TO THE NEXT NANOWORD.

IF A HOLD BIT IS SET ("1") IN THE ACTIVE NANOWORD, THEN THE CORRESPONDING K FIELDS IN THE CONTROL MATRIX DO NOT CHANGE THEIR VALUES AS A RESULT OF GATING THE NEXT NANOWORD INTO THE CONTROL MATRIX. THE ACTION OF THE HOLD BITS IS SUPPRESSED IF THE NEXT NANOWORD IS INVOKED BY MICROINSTRUCTION ENTRY (GATE NS AND LOAD R31 IN THE FINAL T-VECTOR) OR BY PROGRAM CHECK INTERRUPT.

FOR OBVIOUS REASONS, THE HOLD CONTROL AND THE ALLOW INTERRUPT CONTROLS SHOULD NORMALLY BE USED WITH MUTUAL EXCLUSION.

THE FOLLOWING IS A LIST OF K-VECTOR FIELDS AFFECTED BY THE HOLD BITS:

HOLD
=====

KALC
KSHC
KSHA
KS

HOLD 2
=====

KA
KB

4.6 EXTERNAL INTERFACE

THE MATERIAL IN THIS SECTION FUNCTIONALLY SPECIFIES THE INTERFACE BETWEEN THE QM-1 COMPUTER ITSELF AND ITS ENVIRONMENT. INFORMATION ON NANODATA CHANNEL CONTROL UNITS IS PRESENTED IN A SECTION 8 OF THIS MANUAL.

THE EXTERNAL INTERFACE CONSISTS OF EIGHT "EXTERNAL PORTS", EACH IDENTIFIED BY ITS ASSOCIATION WITH A PORT REGISTER (E0 THROUGH E7; SEE SECTION 4.2.5).

THE FOLLOWING "OUTGOING" EXTERNAL INTERFACE FACILITIES ARE BUSSED FOR COMMON USE BY THE EIGHT PORTS:

- A) THE "PHANTOM BUS" (CURRENT INPUT TO THE "PHANTOM" REGISTER, FIPH) SUPPLIES SIX BITS OF INFORMATION
- B) THE "G-BUS" SUPPLIES SIX BITS OF INFORMATION TAKEN FROM ONE OF 16 SOURCES: G0 THROUGH G11, IOID, B, KS, KT. THE SELECTION OF THE SOURCE IS PERFORMED BY THE "GSPEC" FIELD IN THE CURRENTLY ACTIVE T-VECTOR, AND IS FURTHER DISCUSSED IN SECTION 5.5.2 (WHERE THE VALUE ON THE G-BUS IS REFERRED TO AS "G(GSPEC)").
- C) IO CLOCK - A SYNCHRONIZING SIGNAL TO EXTERNAL DEVICES AVAILABLE AT THE PORT DURING EACH T-STEP.
- D) XIO STROBE - A SYNCHRONIZING SIGNAL TO EXTERNAL DEVICES, GENERATED ONLY WHEN XIO IS PRESENT.
- E) MASTER CLEAR - A SIGNAL SENT WHEN THE SYSTEM IS INITIALLY CLEARED. THIS SIGNAL CANNOT BE GENERATED BY PROGRAM CONTROL.

THE FOLLOWING "OUTGOING" EXTERNAL INTERFACE FACILITIES ARE LOCAL TO EACH PORT:

- A) A PATH THROUGH WHICH AN EXTERNAL UNIT CAN READ THE CONTENTS OF THE PORT REGISTER (18 BITS IN PARALLEL).
- B) THE "PORT-XIO" PULSE.
- C) THE "PORT-RIO" PULSE.

THE FOLLOWING "INCOMING" EXTERNAL INTERFACE FACILITIES ARE ALSO LOCAL TO EACH PORT:

- A) A PATH THROUGH WHICH AN EXTERNAL UNIT CAN SUPPLY DATA TO THE PORT REGISTER (18 BITS IN PARALLEL).
- B) THE "IO ID" LINES, THROUGH WHICH AN EXTERNAL UNIT CAN SUPPLY SIX BITS OF INFORMATION TO THE PORT.
- C) SOME NUMBER OF INTERRUPT LEVELS, LOGICALLY ASSIGNED TO THE PORT BY SOFTWARE IN ACCORDANCE WITH THE PHYSICAL SYSTEM CONFIGURATION.

PROGRAM CONTROL OF THE EXTERNAL INTERFACE INVOLVES THESE INTERNAL FACILITIES:

- A) KA (OF THE ACTIVE K-VECTOR),
- B) THE RIO NANOPRIMITIVE ("READ IO"),
- C) THE XIO NANOPRIMITIVE ("TRANSMIT IO"),
- D) SIX-BIT TRANSFER NANOPRIMITIVES, USED TO READ IO ID,
- E) THE INTERRUPT STRUCTURE, AS PRESENTED IN SECTION 4.5.2.4.

KA IS USED MODULO 8 TO SELECT ONE OF THE EIGHT PORTS FOR NANOPRIMITIVE CONTROL. IF NO EXTERNAL UNIT INTERFACES TO THE KA-SELECTED PORT, INCOMING VALUES ARE ZERO AND OUTGOING OPERATIONS ARE NULL.

EXECUTION OF THE XIO NANOPRIMITIVE CAUSES A PORT-XIO SIGNAL TO BE SENT THROUGH THE KA-SELECTED PORT FOR THE DURATION OF THE XIO, AND XIO STROBE TO BE SENT TO ALL PORTS.

EXECUTION OF THE RIO NANOPRIMITIVE HAS THESE EFFECTS:

- A) AT LEADING EDGE, A PORT-RIO PULSE IS SENT THROUGH THE KA-SELECTED PORT.
- B) AT LEADING EDGE, THE KA-SELECTED PORT REGISTER IS SET TO ZEROS.
- C) AT TRAILING EDGE, THE CONTENTS OF THE 18 INCOMING DATA LINES ASSOCIATED WITH THE LATTER REGISTER ARE GATED INTO THAT REGISTER.

THE 6-BIT IO ID OF THE CURRENTLY INTERRUPTING DEVICE IS AVAILABLE TO THE PROGRAM AS AN AUX OR G(GSPEC); SEE SECTION 5.5.2, AND SECTION 8.

4.7 WRITING NANOSTORE

EACH 360-BIT WORD IN NANOSTORE IS PARTITIONED INTO 20 18-BIT BYTES FOR THE PURPOSE OF WRITING. WHEN THE WRITE NS NANOPRIMITIVE IS TO BE USED,

THE NANOSTORE ADDRESS IS TAKEN FROM THE 10 BITS OF R31 ON THE HIGH-ORDER SIDE OF B;

THE BYTE SELECTION IS THE B FIELD, USED MODULO 32;

THE 18 BITS OF DATA TO BE WRITTEN ARE TAKEN FROM THE EOD BUS.

IF EITHER A NONEXISTENT WORD-LOCATION IS ADDRESSED, OR B (MODULO 32) IS GREATER THAN 19 (BYTES ARE ADDRESSED 0 THROUGH 19), WRITE NS IS A NULL OPERATION.

IF WRITE NS AND READ NS APPEAR IN THE SAME ACTIVE T-VECTOR, READ NS IS IGNORED.

4.8 READ-ONLY MEMORIES AND MACHINE START

IN ADDITION TO NANOSTORE AND CONTROL STORE, THE QM-1 CONTAINS A READ-ONLY NANOSTORE (RONS) OF 32 360-BIT WORDS, AND A READ-ONLY CONTROL STORE (ROCS) OF 128 18-BIT WORDS. THESE MEMORIES ARE LOGICALLY DISTINCT FROM NS AND CS, AND ARE ACCESSED AS FOLLOWS:

WHEN THE MEMORY MODE CONTROL BITS IN FIDX ARE CLEARED ("0"), RONS AND ROCS ARE INACCESSIBLE. WHEN THEY ARE SET ("1"), RONS AND ROCS ADDRESS SPACES ARE EFFECTIVELY SUBSTITUTED FOR NS AND CS ADDRESS SPACES, RESPECTIVELY, FOR READ ACCESSES (READ CS, READ NS).

THE CONTENTS OF RONS AND ROCS ARE SPECIFIED BY THE USER AND PERMANENTLY INSERTED BY NANODATA AT INSTALLATION TIME. (NANODATA-SUPPLIED MACHINE DIAGNOSTIC ROUTINES MUST BE INCLUDED, AND NANODATA-SUPPLIED SYSTEM SOFTWARE MAY BE SPECIFIED.)

MACHINE USE OF THESE MEMORIES IS FOR NANOPROGRAM ENTRY AT RONS[0] FOR PROGRAM CHECK (SECTION 4.5.2.2) AND FOR MACHINE START.

WHEN THE QM-1 MASTER CLEAR / START BUTTON IS DEPRESSED,

- A) FIDX IS CLEARED;
- B) THE PROGRAM CHECK STATUS FIELDS (IN E20 AND E21) ARE CLEARED SO THAT THE PROGRAM STARTING AT RONS[0] CAN RECOGNIZE ITS INVOCATION BY MACHINE START RATHER THAN PROGRAM CHECK (SEE SECTION 4.5.2.2); AND
- C) RONS[0] IS FETCHED AND LOADED INTO THE CONTROL MATRIX TO BEGIN EXECUTION.

5 QM-1 FUNCTIONAL SPECIFICATIONS, PART II

5.1 GENERAL

SECTIONS 4 AND 5 OF THIS MANUAL ARE A COMPLETE FUNCTIONAL SPECIFICATION OF THE QM-1 CPU IN TWO PARTS. PART I (SECTION 4) HAS EXPLAINED QM-1 CONCEPTS, ARCHITECTURE AND OPERATIONS. IT HAS PROVIDED AN OVERVIEW OF ALL OF THE FEATURES OF THE MACHINE. PART II (SECTION 5) IS INTENDED AS A PROGRAMMER'S REFERENCE GUIDE AND WILL COMPLETE THE DESCRIPTION OF THOSE PARTS OF THE MACHINE COVERED ONLY BRIEFLY IN PART I.

IT IS ASSUMED THAT THE READER HAS A GENERAL UNDERSTANDING OF THE QM-1 AT THIS POINT. THUS PART II WILL CONCENTRATE MORE ON THE DETAILED OPERATION OF THE INDIVIDUAL MACHINE FUNCTIONS AND LESS ON THEIR POSSIBLE COMBINED USE.

THE NEXT TWO SECTIONS PRESENT, IN SUMMARY FORM, ALL OF THE CONTROL FUNCTIONS INCLUDED IN THE MACHINE. IN MOST CASES, THE FUNCTIONS ARE ACTIVATED BY SINGLE BITS OR CONTAIN A STRING OF BITS USED AS A NUMERIC VALUE. IN THE FEW REMAINING CASES, THE ENCODINGS OF THE BITS ARE GIVEN. REFERENCES ARE PROVIDED FOR EACH FUNCTION TO THE SECTION NUMBERS WHERE THE FUNCTION IS SPECIFIED IN DETAIL.

5.2 SUMMARY OF NANOPRIMITIVE CONTROLS

5.2.1 K-VECTOR CONTROL FIELDS

THE CONTROL FUNCTION OF EACH OF THE FIELDS IN THE K-VECTOR IS SUMMARIZED IN THE TABLE BELOW, ALONG WITH REFERENCES TO SECTIONS IN WHICH THE FUNCTION IS DESCRIBED. (THE NUMBER OF BITS IN THE FIELD IS SHOWN IN PARENTHESES.)

CONTROL FIELD =====	(BITS) =====	SUMMARY OF CONTROL FUNCTION =====	REFERENCES =====
KN	(10)	ADDRESS OF POSSIBLE SUCCESSOR NANOWORD. NANOBRANCH ADDRESS AND SOURCE FOR NPC LOAD.	4. 5. 2. 3
SUPERVISOR	(1)	PROGRAM CHECK IF ON WHEN THIS WORD IS INVOKED WHILE NOT IN SUPERVISOR MODE.	RESERVED- FUTURE USE
LEGAL MICRO ENTRY	(1)	PROGRAM CHECK IF NOT ON WHEN THIS WORD IS INVOKED BY A MICROINSTRUCTION.	4. 5. 3
BRANCH	(1)	MUST BE ON IF NANOBRANCH PLANNED FROM THIS WORD. COMPLEMENTED AFTER EACH READ NS WHEN ALTERNATE IS ON.	4. 5. 2. 3
ALTERNATE	(1)	CAUSES BRANCH TO BE COMPLEMENTED AFTER EACH READ NS.	4. 5. 2. 3
HOLD	(1)	INHIBITS AUTOMATIC LOADING OF KALC, KSHC, KSHA, AND KS FROM NEXT NANOWORD TO BE EXECUTED, UNLESS EXECUTED BY MICROINSTRUCTION OR PROGRAM CHECK.	4. 5. 2. 4 4. 5. 4. 2
HOLD 2	(1)	INHIBITS AUTOMATIC LOADING OF KA AND KB FROM NEXT NANOWORD TO BE EXECUTED, UNLESS EXECUTED BY MICROINSTRUCTION OR PROGRAM CHECK.	4. 5. 2. 4 4. 5. 4. 2
ALLOW NANO INTERRUPT	(1)	ALLOWS HIGHER-PRIORITY INTERRUPTS AT END OF EXECUTION OF THIS WORD, IF NANOBRANCH IS NOT TAKEN.	4. 5. 2. 4 4. 5. 4. 2

ALLOW MICRO INTERRUPT	(1)	ALLOWS LOWER-PRIORITY INTERRUPTS AT END OF EXECUTION OF THIS WORD, IF NANOBANCH IS NOT TAKEN.	4.5.2.4 4.5.4.2
GENERATE INTERRUPT	(1)	GENERATES OR CLEARS AN INTERRUPT LEVEL ACCORDING TO GIGSPECJ IN T1.	4.5.2.4 5.8.1
ALU STATUS ENABLE	(1)	ENABLES MOVE OF C, S, R, O BITS FROM LOCAL TO GLOBAL UPON GATE ALU; C TREATED SPECIALLY.	5.6.2
SH STATUS ENABLE	(1)	ENABLES MOVE OF SHB, SLB BITS FROM LOCAL TO GLOBAL UPON GATE SH.	5.6.3
DIRECT MS ACCESS	(1)	INHIBITS MS BASE ADDRESSING AND FIELD LENGTH PROTECTION IN THIS NANOWORD.	4.2.6.3
KA	(6)	CONSTANT AND/OR SCRATCH FIELD FOR NANOWORD; SOURCE AND DESTINATION AUX.	4.3.3
KB	(6)	CONSTANT AND/OR SCRATCH FIELD FOR NANOWORD; SOURCE AND DESTINATION AUX.	4.3.3
KALC	(6)	ALU CONTROL; DESTINATION AUX.	5.6.2
KSHC	(6)	SHIFT CONTROL; DESTINATION AUX.	5.6.3
KSHA	(6)	SHIFT AMOUNT; DESTINATION AUX.	5.6.3
KS	(6)	GLOBAL CONDITION (AND GENERAL) TEST MASK; SOURCE AND DESTINATION AUX.	4.6 5.7.1
KT	(6)	LOCAL CONDITION TEST MASK (ALSO CONSTANT AND/OR SCRATCH); SOURCE AND DESTINATION AUX.	5.7.1
KX	(6)	SPECIAL CONDITION TEST MASK (ALSO CONSTANT AND/OR SCRATCH); SOURCE AND DESTINATION AUX.	5.7.1
SPARE	(2)	RESERVED FOR FUTURE USE	

5.2.2 T-VECTOR CONTROL FIELDS

THE CONTROL FUNCTION OF EACH OF THE FIELDS IN THE ACTIVE T-VECTOR IS SUMMARIZED IN THE TABLE BELOW, ALONG WITH REFERENCES TO SECTIONS IN WHICH THE FUNCTION IS DESCRIBED. A CODE SHOWING THE CHARACTERISTIC TIMING OF THE ACTION ASSOCIATED WITH THE FUNCTION IS GIVEN; LE = LEADING EDGE, TE = TRAILING EDGE. (THE NUMBER OF BITS IN THE FIELD IS SHOWN PARENTHETICALLY.)

CONTROL FIELD =====	BITS =====	SUMMARY OF CONTROL FUNCTION =====	TIME =====	REFS. =====
STRETCH	(1)	STRETCHES TIME OF THIS T-STEP FROM ONE T-PERIOD TO TWO.		4.4
WRITE NS	(1)	WRITES 18 BITS FROM EOD BUS INTO NANOSTORE	LE	4.7 5.4.1.2
XIO	(1)	SENDS PULSE TO EXTERNAL INTERFACE; ONE OF EIGHT EXTERNAL PORTS SELECTED BY KA.	LE	4.6
RIO	(1)	CLEARs PORT REGISTER AND SENDS PULSE THROUGH PORT, THEN GATES EXTERNAL DATA WORD INTO PORT REGISTER; SELECTED BY KA.	LE	4.6
MSGO	(1)	INITIATES MS OPERATION; SPLIT-CYCLE IF ALONE, FULL-READ IF MSRS SIMULTANEOUS.	LE	4.2.6.2 5.4.3
MSRS	(1)	IF ALONE, REQUESTS SECOND HALF-CYCLE OF MS SPLIT-CYCLE OPERATION; IF WITH MSGO, INITIATES FULL-READ.	LE	4.2.6.2 5.4.3
GATE MS	(1)	GATES MOD BUS INTO LOCAL STORE OR PORT REGISTERS; MODIFIED BY RMI SELECT.	TE	4.2.6.1 5.4.3
RMI SELECT 00 BYPASS 01 PARAMETER SET A 10 PARAMETER SET B 11 PARAMETER SET C	(2)	SELECTS RMI PARAMETERS FOR GATE MS, INCLUDING BYPASS. IF RMI NOT INSTALLED ALL ENCODINGS ARE BYPASS	LE	4.2.6.4
GATE ES	(1)	GATES EOD BUS INTO LOCAL STORE.	TE	4.2.5

LOAD ES	(1)	LOADS AN EXTERNAL STORE REGISTER FROM EID BUS.	TE	4. 2. 5
READ CS	(1)	READS CONTROL STORE; USES CS ADDR SELECT.	LE	4. 2. 4 5. 4. 2. 2
WRITE CS	(1)	WRITES CONTROL STORE; USES CS ADDR SELECT.	LE	4. 2. 4 5. 4. 2. 3
CS ADDR SELECT	(3)	SELECTS ADDRESS FOR READ CS, WRITE CS. (MPC IS SELECTED BY FMPC) A AND AB ARE SIGN EXTENDED OPERANDS. SWITCHES ARE MANUALLY SET IN MACHINE.	LE	4. 2. 4 5. 4. 2. 1
000 CIA				
001 COD				
010 MPC				
011 MPC+1				
100 MPC+2				
101 MPC+B				
110 MPC+AB				
111 SWITCHES				
GATE CS	(1)	GATES COD BUS INTO LOCAL STORE.	TE	4. 2. 4 5. 4. 2
GATE ALU	(1)	GATES ADD BUS INTO LOCAL STORE.	TE	4. 2. 3
GATE SH	(1)	GATES SOD BUS INTO LOCAL STORE.	TE	4. 2. 3
CARRY CTL	(3)	CONTROLS CARRY OPERATION WITHIN THE ALU AND SHIFTER COMPONENTS.	TE	4. 2. 3. 4
000 NO OPERATION				
001 CLEAR CIH				
010 SET CIH				
011 ALU TO BOTH				
100 ALU TO COH				
101 SET COH				
110 CLEAR COH				
111 SH TO COH				
INDEX	(1)	GATES INDEX ALU OUTPUT INTO LOCAL STORE, SELECTED BY G(GSPEC).	TE	4. 2. 2. 3 5. 6. 4
INC MPC	(1)	INCREMENTS MPC SELECTED BY FMPC; MODIFIED BY GSPEC.	TE	4. 2. 2 5. 4. 2. 1 5. 6. 5

LOAD NPC	(2)	LOADS OR SEQUENCES NANOPROGRAM COUNTER.	TE	4.5.3
00 NO OPERATION				4.5.4
01 (CS)				
10 (KN)				
11 (SEQUENCE)				
READ NS	(1)	READS NS; ADDRESS IS FROM PRIORITY-SELECT MECHANISM. INFLUENCES BRANCH.	LE	4.5 5.4.1.1
GATE NS UNCONDITIONALLY	(1)	CAUSES THE NANOWORD LAST READ TO BE GATED INTO THE CONTROL MATRIX. INDEPENDENT OF ANY TEST ACTION IN T.	TE	4.5.1 5.5.1
TEST ACTION	(1)	CONDITIONAL ACTION BASED ON TEST SPECIFIER	TE	4.5 4.5.1 5.7.2
0 SKIP				
1 GATE NS				
TEST SPECIFIER	(3)	SPECIFIES THE CONDITIONS UNDER WHICH TEST ACTION IS TO BE EXECUTED	LE	5.7.1
000 NEVER				
001 ALWAYS				
010 IF FIST AND KS = 0				
011 IF FIST AND KS NOT = 0				
100 IF LOCAL CONDS AND KT = 0				
101 IF LOCAL CONDS AND KT NOT = 0				
110 IF SPECIAL CONDS AND KX = 0				
111 IF SPECIAL CONDS AND KX NOT = 0				
LOAD R31	(1)	ENABLES R31 TO BE LOADED WITH MICRO-INSTRUCTION PARAMETERS.	TE	4.5.3.2 5.3.4
AUXILIARY ACTION	(1)	INITIATES ACTION SPECIFIED BY THE CONTENTS OF FACT (F REGISTER 14).	LE	4.3.2.3 5.8.2
GSPEC	(4)	SELECTS A G OR PSEUDO-G FOR 6-BIT TRANSFERS, RIGHT INPUT TO ALUF, USED IN GENERATE INTERRUPT, EXTERNAL INTERFACE G-LINES, ALSO USED WITH INC NPC.		5.5.2
0000 G0				

1011 G11				
1100 IO IO				
1101 B				
1110 KS				
1111 KT				

FSEL0	(5)	SELECTS F REGISTER FOR 6-BIT TRANSFERS	5.5.2
FSEL1	(5)	IN GROUP 0, 1, AND 2 RESPECTIVELY.	
FSEL2	(5)		
AUX0	(3)	SELECTS AUX FOR 6-BIT TRANSFERS IN	
AUX1	(3)	GROUP 0, 1, AND 2 RESPECTIVELY.	
AUX2	(3)	(AUX2 APPLIES TO GROUP 2 INPUT,	
AUX3	(3)	AUX3 APPLIES TO GROUP 2 OUTPUT.)	
IN0	(1)	COMMANDS AUX INTO F REGISTER TRANSFER	
IN1	(1)	USING AUX0, AUX1, AUX2 TO FSEL0,	
IN2	(1)	FSEL1, FSEL2 RESPECTIVELY.	
OUT1	(1)	COMMANDS F REGISTER OUTPUT TO AUX	
OUT2	(1)	TRANSFER USING FSEL0, FSEL1, FSEL2 TO	
OUT3	(1)	AUX0, AUX1, AUX3 RESPECTIVELY.	
SPARE	(1)		


72 BITS

5.3 FUNDAMENTAL TIMING CONSIDERATIONS

5.3.1 GENERAL

THE "HARDWARE LEVEL" QM-1 IS A HIGHLY PARALLEL MACHINE. ONE OF THE TASKS FACING THE NANOPROGRAMMER IS TO PUT TOGETHER THE FUNCTIONS HE DESIRES IN SUCH A WAY AS TO UTILIZE THIS PARALLELISM TO THE FULLEST EXTENT POSSIBLE. HENCE HE MUST HAVE AN INTIMATE KNOWLEDGE OF THE INTERNAL TIMING OF THE MACHINE. THIS SECTION ON TIMING CONSIDERATIONS IS INCLUDED IN ORDER THAT NANOPROGRAMMERS CAN ANSWER QUESTIONS REGARDING MEANINGFUL COMBINATIONS OF FUNCTIONS IN THE SAME OR ADJACENT T-STEPS.

ALL T-VECTOR CONTROL FUNCTIONS HAVE BEEN CLASSIFIED AS "LEADING EDGE" (LE) OR "TRAILING EDGE" (TE) FUNCTIONS DEPENDING ON THE TIME OF THE ACTION THEY INITIATE, RELATIVE TO THE PERIOD OF THE T-STEP IN WHICH THE CONTROL IS ACTIVE. THE PERIOD OF ANY T-STEP IS DEFINED AS THE TIME BETWEEN THE MACHINE CLOCK PULSE WHICH CAUSES THE T-VECTOR TO BECOME ACTIVE AND THE NEXT CLOCK PULSE WHICH CAUSES THE NEXT T-VECTOR TO BECOME ACTIVE. THESE PULSES ARE KNOWN AS T-CLOCK PULSES OR JUST T-CLOCKS AND THE T-STEP (T2 FOR EXAMPLE) IS AS SHOWN.

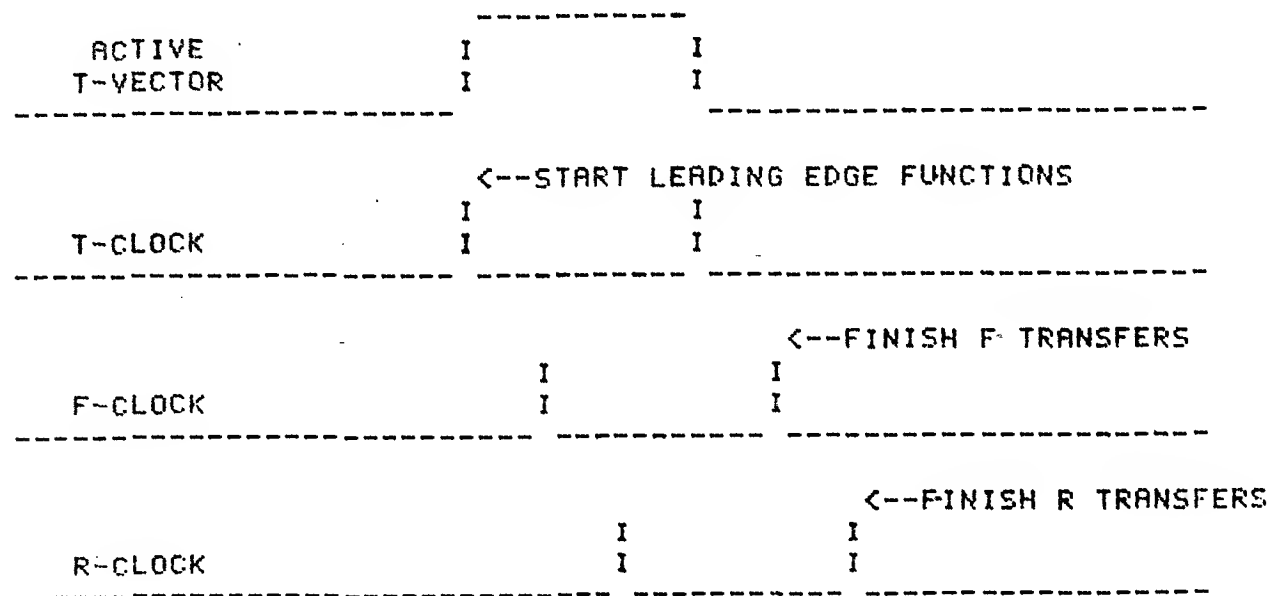


LEADING EDGE FUNCTIONS ARE THOSE WHICH ARE TRIGGERED BY THE BEGINNING EDGE (OR ACTIVATION) OF THE T-STEP AND TRAILING EDGE FUNCTIONS ARE THOSE TRIGGERED BY THE ENDING EDGE (OR DEACTIVATION) OF THE T-STEP.

EXAMINATION OF THIS SITUATION FOR TWO SUCCESSIVE T-STEPS SHOWS THAT A TRAILING EDGE FOR ONE T-STEP OCCURS AT EXACTLY THE SAME TIME AS THE LEADING EDGE OF THE NEXT. THUS IT WOULD SEEM THAT A LEADING EDGE EVENT COULD OCCUR AT EXACTLY THE SAME TIME A TRAILING EDGE TRANSITION IS HAPPENING. IN ACTUALITY, THIS PROBLEM IS AVOIDED BY HAVING SOME FUNCTIONS "MORE TRAILING EDGE" THAN OTHERS. THIS IS NECESSARY SINCE THE MACHINE STATE VECTOR (ACTIVE K AND T-VECTORS) MUST BE IN A DEFINED STATE BEFORE THE 6-BIT DOMAIN CAN OPERATE PROPERLY. AND THE 6-BIT SECTIONS MUST BE IN A DEFINED STATE FOR THE 18 BIT DOMAIN TO OPERATE PROPERLY. HENCE THE NECESSITY OF TWO ADDITIONAL CLOCKS DERIVED FROM THE T-CLOCK. THESE ARE, NOT SUPRISINGLY, CALLED THE "F REGISTER CLOCK" (OR F-CLOCK) AND THE "LS REGISTER CLOCK" (OR R-CLOCK). THE ACTUAL DELAYS BETWEEN THESE CLOCKS ARE

IMPORTANT ONLY WHEN THE BOUNDARIES BETWEEN THE THREE DOMAINS WITHIN THE MACHINE ARE CROSSED. THE EXTREME CASE IS IN R31 OPERATIONS SINCE ALL THREE DOMAINS MEET IN R31 (COVERED EXTENSIVELY IN SECTION 5.3.5).

BRIEFLY, THE CLOCKS ACT AS FOLLOWS. THE T-CLOCKS ACTIVATE A T-VECTOR. ALL LEADING EDGE FUNCTIONS ARE BEGUN IMMEDIATELY. ALL DECODING AND SET-UP FOR TRAILING EDGE FUNCTIONS ALSO BEGINS AT THIS TIME. THE NEXT T-CLOCK DEACTIVATES THIS T-VECTOR (WHICH WILL NOW BE CALLED THE PREVIOUS T-VECTOR) AND ACTIVATES THE NEXT ONE. APPROXIMATELY 20 NANOSECONDS AFTER THE T-CLOCK, THE F-CLOCK OCCURS, COMPLETING ANY 6-BIT DATA TRANSFERS SPECIFIED IN THE PREVIOUS T-VECTOR. APPROXIMATELY 20 NANOSECONDS AFTER THE F-CLOCK, THE R-CLOCK OCCURS. THIS COMPLETES ANY 18-BIT TRANSFERS SPECIFIED IN THE PREVIOUS T-VECTOR. THIS SEQUENCE IS SHOWN IN FIGURE 5.3.1A.



RELATIONSHIP OF MACHINE CLOCKS

FIGURE 5.3.1A

5.3.2 LEADING EDGE FUNCTIONS

THE MEMORY REFERENCE FUNCTIONS; READ/WRITE NANOSTORE, READ/WRITE CONTROL STORE, GO/RESTART MAIN STORE, ARE ALL OPERATIONS WHICH DO NOT CAUSE DATA TO BE GATED BUT ARE NECESSARY TO MAKE DATA AVAILABLE FOR GATING. IN ORDER TO MAKE THE DATA AVAILABLE AS SOON AS POSSIBLE, THESE OPERATIONS MUST BE INITIATED AS SOON AS THE T-VECTOR IN WHICH THEY ARE SPECIFIED BECOMES ACTIVE. THUS THEY ARE LEADING EDGE FUNCTIONS. CARE MUST BE TAKEN TO ASSURE THAT THE ADDRESS AND/OR DATA TO BE USED BY THE OPERATION IS STABLE BEFORE THE FUNCTION IS INITIATED. SECTION 5.4 COVERS THIS IN DETAIL.

XIO AND RIO ARE THE ONLY OTHER LEADING EDGE FUNCTIONS. BOTH XIO AND RIO GENERATE A SIGNAL TO THE EXTERNAL PORT AND MUST THEREFORE BEGIN ON THE LEADING EDGE IN ORDER FOR THE ACTION THEY INITIATE TO BE COMPLETED BY THE END OF THE T-PERIOD.

5.3.3 TRAILING EDGE EVENTS

5.3.3.1 T-CLOCK EVENTS

LOAD NPC IS EXECUTED ON THE TRAILING EDGE T-CLOCK. THE THREE POSSIBLE SOURCES OF NEW VALUES TO BE TRANSFERRED INTO THE NANOPROGRAM COUNTER ARE THE COD BUS, KN IN THE EXECUTING NANOWORD, AND THE CURRENT VALUE OF NPC. NO SPECIAL TIMING PROBLEMS ARISE WITH THE LOAD NPC OPERATION.

GATE NS AND LOAD R31 ARE THE ONLY OTHER FUNCTIONS EXECUTED ON THE TRAILING EDGE T-CLOCK. BOTH ARE USED TO CAUSE THE TRANSITION BETWEEN ONE INSTRUCTION AND THE NEXT. THUS THEY MUST BE COMPLETED BEFORE ANY OTHER FUNCTIONS CAN BEGIN.

A CONFLICT IS POSSIBLE BETWEEN EACH OF THESE FUNCTIONS AND SOME OTHER FUNCTION IN THE MACHINE. IN THE CASE OF GATE NS, AN F-TRANSFER INTO A K, INITIATED IN THE SAME T-STEP AS THE GATE NS WILL CAUSE AN UNDEFINED RESULT IN THE K UNLESS THE K IS "HELD" BY THE APPROPRIATE HOLD COMMAND. SEE SECTION 4.5.4.3.

IN THE CASE OF LOAD R31, AN F TRANSFER INTO R31, INITIATED IN THE SAME T-STEP AS LOAD R31 WILL CAUSE AN UNDEFINED RESULT IN R31. SEE SECTION 5.3.4. AN R TRANSFER WILL OVERRIDE THE EFFECT OF THE LOAD R31.

5.3.3.2 F-CLOCK EVENTS

ALL F TRANSFERS ARE COMPLETED BY THE TRAILING EDGE F-CLOCK. THESE INCLUDE F REGISTER INCREMENT AND DECREMENT AND ALUF OPERATIONS SINCE THE RESULTS ARE GATED AS AN F TRANSFER.

SIMULTANEOUS F TRANSFERS TO THE SAME DESTINATION DO NOT CAUSE UNDEFINED RESULTS SINCE THEY OCCUR AT EXACTLY THE SAME TIME. A LOGICAL "OR" OF THE TRANSFERRED VALUES OCCURS.

5.3.3.3 R-CLOCK EVENTS

ALL 18 BIT TRANSFERS INTO LOCAL STORE OR EXTERNAL STORE ARE SYNCHRONIZED ON THE TRAILING EDGE R-CLOCK. SIMULTANEOUS R TRANSFERS INTO THE SAME LOCAL STORE REGISTER WILL PRODUCE THE LOGICAL "OR" OF THE TRANSFERRED VALUES.

5.3.4 R31 OPERATIONS

R31 IS THE PRIMARY INTERFACE BETWEEN THE 18-BIT ARCHITECTURE AND THE 6-BIT ARCHITECTURE IN THE QM-1, SINCE IT IS SIMULTANEOUSLY A LOCAL STORE REGISTER AND ALSO CONTAINS THREE 6-BIT AUX FIELDS:

```

-----
R31  I          18 BIT LOCAL STORE REGISTER          I
-----

```

```

-----
R31  I          C          I          A          I          B          I
-----
BIT  17 16 15 14 13 12 11 10 09 08 07 06 05 04 03 02 01 00

```

WHEN SERVING IN ITS ADDITIONAL SPECIAL FUNCTION AS THE MICRO INSTRUCTION REGISTER, R31 IS CLASSIFIED IN THE CONTROL MATRIX DOMAIN, SINCE THE LOAD R31 COMMAND TRANSFERS THE SAVED MICROINSTRUCTION PARAMETER PART INTO R31 ON THE LEADING EDGE T-CLOCK:

```

-----
R31  I          7 ZERO BITS          I          A          I          B          I
-----

```

BECAUSE OF THIS INTERFACE, CARE MUST BE TAKEN IN ORGANIZING TRANSFERS INVOLVING R31. THE FOLLOWING PROGRAMMING RULES ARE DERIVED FROM THE CLOCK RELATIONSHIPS DISCUSSED IN SECTION 5.3.1.

1. IF THE ONLY TRANSFERS INTO R31 COMMANDED IN A T-STEP ARE IN THE 18 BIT DOMAIN, THEN SIX-BIT TRANSFERS OUT OF R31 COMMANDED IN THE FOLLOWING T-STEP WILL OCCUR TOO SOON TO USE THE 18-BIT VALUE, AND WILL INSTEAD TRANSFER AN UNDEFINED VALUE. THIS HOLDS TRUE UNLESS THE T-STEP CONTAINING THE SIX BIT TRANSFER COMMAND IS STRETCHED, IN WHICH CASE SUCH TRANSFERS DO USE THE NEW 18-BIT VALUE.
2. IF THE ONLY TRANSFERS INTO R31 COMMANDED IN A T-STEP ARE FROM THE SIX-BIT DOMAIN, THESE VALUES ARE AVAILABLE FOR TRANSFERRING OUT TO EITHER DOMAIN IN THE NEXT T-STEP.

3. IF TRANSFERS INTO R31 FROM BOTH THE SIX-BIT AND THE 18-BIT DOMAINS ARE COMMANDED IN A T-STEP, AN UNDEFINED VALUE RESULTS IN R31.
4. WHEN R31 IS USED WITHIN ONE DOMAIN ONLY, NORMAL TIMING RULES APPLY, AS IN SECTION 4.4.
5. IF SIX-BIT TRANSFERS INTO R31 ARE CONCURRENT WITH LOAD R31 PARAMETER LOADING (I.E., IF SUCH TRANSFERS ARE COMMANDED IN THE LAST T-STEP OF A NANOWORD ALONG WITH THE LOAD R31 COMMAND) AN UNDEFINED VALUE RESULTS IN R31.

5.4 MEMORY REFERENCE FUNCTIONS

5.4.1 NANOSTORE OPERATIONS

5.4.1.1 READ NS

READ NS IS A LEADING EDGE COMMAND THAT CAUSES A 360-BIT WORD TO BE READ FROM NANOSTORE USING THE NANOSTORE ADDRESS SELECTED PREVIOUSLY BY THE NANOSTORE ADDRESSING MECHANISM (SECTION 4.5.2). THE ADDRESS MUST HAVE BEEN ESTABLISHED IN THE PREVIOUS T-STEP. THIS ADDRESS MUST BE STABLE FOR THE T-STEP IN WHICH READ NS OCCURS. NEITHER OF THESE REQUIREMENTS CAUSE ANY DIFFICULTY SINCE THE LOAD NPC COMMANDS ARE PROPERLY SYNCHRONIZED TO SATISFY THEM.

FOR NANOSTORE DATA-OUT TO BE AVAILABLE IN TIME FOR A TRAILING-EDGE GATE NS EXECUTED IN A GIVEN T-PERIOD, THE LEADING-EDGE OPERATION READ NS MUST OCCUR IN THE PRIOR T-PERIOD OR EARLIER. THEREFORE READ NS, GATE NS CAN BE PROGRAMMED AS A SEQUENCE IN ONE T-STEP IF AND ONLY IF THAT T-STEP IS STRETCHED.

READ NS, GATE NS EXECUTED IN THE SAME UN-STRETCHED T-STEP RESULT IN AN UNDEFINED VALUE LOADED INTO THE CONTROL MATRIX.

READ NS CANNOT BE COMMANDED IN T1 OF ANY NANOWORD THAT ALLOWS INTERRUPTS. UNDEFINED DATA RESULTS IN THIS SITUATION.

5.4.1.2 WRITE NS

WRITE NS IS A LEADING-EDGE NANOPRIMITIVE THAT INITIATES WRITING 18 BITS INTO NANOSTORE FROM THE EOD BUS AT THE ADDRESS SPECIFIED BY THE CONTENTS OF R31. THE B FIELD OF R31 SPECIFIES WHICH OF THE 20 BYTES OF THE NANOWORD IS TO BE WRITTEN (0-19, MODULO 32). THE 10 BITS OF R31 ON THE HIGH ORDER SIDE OF THE B FIELD IN R31 SELECT THE PARTICULAR NANOWORD TO BE WRITTEN. AN INVALID NANOSTORE ADDRESS (OUT OF RANGE) CAUSES THE WRITE NS TO BE IGNORED.

```

-----
R31  IXX XX <-----NANOWORD ADDRESS-----> XX  BYTE ADDRESS I  BITS MARKED XX
-----
      17 16 15 14 13 12 11 10 09 08 07 06 05 04 03 02 01 00  ARE IGNORED.

```

THE WRITE NS OPERATION CONSUMES THREE T-PERIODS. (ATTEMPTS TO INITIATE A READ NS OR ANOTHER WRITE NS DURING SUCH TIME ARE IGNORED.) THUS IF WRITE NS APPEARS IN T-PERIOD T(N), NANOSTORE WILL BE WRITTEN IN TIME FOR A READ NS EXECUTED IN T-PERIOD T(N+3).

ASSUMING THAT WRITE NS IS EXECUTED IN T-PERIOD T(N), RESULTS OF THE OPERATION ARE UNDEFINED IF:

1. THE NANOSTORE WORD-ADDRESS IS MODIFIED BY A COMMAND APPEARING IN T-PERIOD T(N-1).
2. THE BYTE-SELECTOR IN B IS MODIFIED BY A COMMAND APPEARING IN T-PERIOD T(N-1).
3. THE DATA ON THE EOD BUS IS MODIFIED BY A COMMAND APPEARING IN T-PERIOD T(N-1).

5.4.2 CONTROL STORE OPERATIONS

5.4.2.1 CONTROL STORE ADDRESS SELECTION

FOR EITHER READ CS OR WRITE CS, THE ADDRESS USED IS DETERMINED BY THE VALUE OF THE CS ADDR SELECT FIELD IN THE SAME T-VECTOR AS THE READ CS OR WRITE CS. SINCE BOTH READ CS AND WRITE CS ARE LEADING EDGE FUNCTIONS, THE ADDRESS MUST BE STABLE AT THE BEGINNING OF THE T-STEP IN WHICH THE COMMAND OCCURS. THIS HAS DIFFERENT IMPLICATIONS, DEPENDING ON THE ADDRESS SOURCE. EACH OF THE CASES IS COVERED BELOW, ASSUMING THAT THE READ CS OR WRITE CS OCCURS IN T(N):

CS ADDR SELECT			CONDITIONS
0	000	CIA	ADDRESS IS TAKEN FROM THE LOCAL STORE REGISTER DESIGNATED BY FCIA. NO COMMANDS CHANGING FCIA OR CIA SHOULD APPEAR IN T(N-1)
1	001	COD	HERE THE ADDRESS IS TAKEN DIRECTLY FROM THE COD BUS RATHER THAN FROM A REGISTER. SINCE THE ONLY THING THAT CAN CHANGE THE DATA ON THE COD BUS IS A PREVIOUS READ CS, THIS SHOULD NOT OCCUR IN T(N-1) UNLESS STRETCHED.
2	010	MPC	ADDRESS IS TAKEN FROM THE LOCAL STORE REGISTER DESIGNATED BY FMPC. NO COMMANDS CHANGING FMPC OR MPC SHOULD APPEAR T(N-1).
3	011	MPC+1	HERE THE ADDRESS DEPENDS ON FMPC, MPC AND THE OUTPUT OF THE MPC INCREMENT FACILITY. AGAIN, NOTHING THAT CHANGES EITHER FMPC OR MPC SHOULD OCCUR IN T(N-1).
4	100	MPC+2	
5	101	MPC+B	IN THIS CASE, AN ADDED FACTOR IS INVOLVED - THE CONTENTS OF R31. THUS NOTHING THAT CHANGES THE CONTENTS OF R31 SHOULD APPEAR IN T(N-1)
6	110	MBP+AB	
7	111	SWITCHES 12 EXTERNAL SWITCHES DETERMINE THE CONTROL STORE ADDRESS.	

THE CONTROL STORE ADDRESS DETERMINED BY THE ABOVE SELECTION MUST BE STABLE ONLY FOR THE DURATION OF THE T-STEP IN WHICH READ CS OR WRITE CS OCCURS. THUS IT IS POSSIBLE TO SPECIFY IN THE SAME T-STEP, ANY TRAILING EDGE OPERATIONS THAT CHANGE THE ADDRESS.

CONTROL STORE ADDRESSES ARE TREATED MODULO 32K (15 BITS). ATTEMPTS TO READ NONEXISTANT LOCATIONS RESULT IN ZEROS ON THE COD BUS. ATTEMPTS TO WRITE NONEXISTANT LOCATIONS ARE IGNORED.

IF READ CS AND WRITE CS OCCUR SIMULTANEOUSLY, ONLY THE WRITE CS OCCURS. READ CS AND/OR WRITE CS COMMANDS ARE VALID IN TWO SUCCESSIVE T-STEPS ONLY IS THE FIRST T-STEP IS STRETCHED.

5.4.2.2 READ CS

READ CS READS THE 18 BIT VALUE FROM CONTROL STORE AT THE ADDRESS DETERMINED BY CS ADDR SELECT. THE VALUE READ IS PLACED ON THE COD BUS FOR GATING INTO A LOCAL STORE REGISTER, FOR LOADING AS A MICROINSTRUCTION, FOR USE AS AN INDIRECT CONTROL STORE ADDRESS OR FOR USE AS AN ARITHMETIC OPERAND. IF THE READ CS IS EXECUTED IN T-PERIOD $T(N)$, THE DESIRED DATA IS AVAILABLE ON COD FOR EACH OF THESE USES AT THE TRAILING EDGE OF $T(N+1)$. IT IS AVAILABLE AT THE TRAILING EDGE OF $T(N)$ IF THE T-STEP IS STRETCHED.

ONCE ESTABLISHED, A COD VALUE REMAINS AVAILABLE UNTIL CHANGED BY THE NEXT READ CS OPERATION.

5.4.2.2 WRITE CS

WRITE CS WRITES THE 18 BIT VALUE FROM THE LOCAL STORE REGISTER SPECIFIED BY FCID INTO CONTROL STORE AT THE ADDRESS DETERMINED BY CS ADDR SELECT.

IN ADDITION TO ALL OF THE RULES THAT APPLY TO THE ADDRESS SELECTION (SECTION 5.4.2.1), THE DATA ON THE CID BUS MUST BE STABLE FOR THE DURATION OF THE T-STEP IN WHICH THE WRITE CS OCCURS. THUS, FOR A WRITE CS IN $T(N)$ TO BE VALID, NOTHING THAT CHANGES EITHER FCID OR CID MAY APPEAR IN $T(N-1)$.

5.4.3 MAIN STORE OPERATIONS

5.4.3.1 READ MAIN STORE

A READ MAIN STORE OPERATION IS INITIATED BY PLACING THE APPROPRIATE VALUES IN THE FOLLOWING REGISTERS AND SIMULTANEOUSLY ISSUING THE COMMANDS MSGO AND MSRS.

FMIX - POINTER TO LOCAL STORE REGISTER USED FOR MAIN STORE ADDRESS
MIX - MAIN STORE ADDRESS IN LOCAL STORE REGISTER
E16 - MAIN STORE BASE ADDRESS IN EXTERNAL STORE REGISTER
E17 - MAIN STORE FIELD LENGTH IN EXTERNAL STORE REGISTER

THESE FOUR VALUES MUST BE CHANGED NO LATER THAN THE T-STEP BEFORE THE COMMANDS MSGO AND MSRS APPEAR. THE VALUES NEED ONLY BE HELD STABLE DURING THE T-STEP IN WHICH MSGO AND MSRS ARE ISSUED. MSGO AND MSRS MUST NOT BE ISSUED TOGETHER UNTIL A TEST OF "MS BUSY" INDICATES THAT MAIN STORE IS AVAILABLE.

THE WORD IN MAIN STORE, ADDRESSED BY THE SUM OF THE VALUES IN MIX AND E16 IS READ AND PLACED ON THE MOD BUS IN TIME FOR GATING WITH THE TRAILING EDGE COMMAND GATE MS IN THE SAME T-STEP THAT THE TEST "MS DATA" INDICATES DATA AVAILABLE. IF THE VALUE IN MIX EXCEEDS THE VALUE IN E17, OR IF THE ADDRESSED WORD IS BEYOND THE RANGE OF INSTALLED MAIN STORE ADDRESSES, THE RESULT ON THE MOD BUS IS ZERO AND A PROGRAM CHECK OCCURS.

IF DESIRED, THE PARTICIPATION OF E16 AND E17 CAN BE BYPASSED BY ISSUING THE COMMAND DIRECT MS IN THE K-VECTOR OF THE WORD CONTAINING MSGO AND MSRS. IN EFFECT, THIS CAUSES THE OPERATION PERFORMED TO BE EQUIVALENT TO HAVING ZERO IN E16 AND ALL ONES IN E17. ADDRESSING MAIN STORE BEYOND INSTALLED ADDRESSES RESULTS IN A PROGRAM CHECK WITH ALL ONES ON THE MOD BUS.

IF THE MS ADDRESSING AND PROTECTION OPTION IS NOT INSTALLED, THE EFFECT IS THAT OF HAVING DIRECT MS ALWAYS ON. THUS E16 AND E17 NEVER PARTICIPATE IN MAIN STORE ADDRESSING AND ARE AVAILABLE FOR OTHER USE.

THE VALUE IN FMOD MUST BE ESTABLISHED NO LATER THAN THE T-PERIOD IN WHICH THE GATE MS OCCURS. BOTH FMIX AND FMOD HAVE AN EXTENDED ADDRESS FUNCTION AS SHOWN:

VALUE	FMIX	FMOD
0 - 31	LOCAL STORE REGISTERS	LOCAL STORE REGISTERS
32 - 39	EXTERNAL STORE E0 - E7	EXTERNAL STORE E0 - E7
40 - 64	SOURCE OF ALL ONES	NULL OPERATION

THIS PERMITS BOTH ADDRESS AND DATA FOR MAIN STORE OPERATIONS TO BE PLACED IN THE EXTERNAL STORE PORT REGISTERS AS WELL AS IN LOCAL STORE.

AFTER READ MAIN STORE HAS BEEN INITIATED, ANY SUBSEQUENT MSGO OR MSRS SIGNALS WILL BE IGNORED UNTIL MS BUSY IS TURNED OFF AT THE END OF THE FULL MEMORY CYCLE. THE MINIMUM TIMING BETWEEN LEADING EDGE OF MSGO AND TRAILING EDGE OF GATE MS IS 640 NANOSECONDS. THE MINIMUM PERIOD OF SUCCESSIVE MSGO SIGNALS IS 800 NANO-SECONDS.

5.4.3.2 WRITE MAIN STORE

SINCE THE MIX BUS IS SHARED FOR BOTH ADDRESS AND DATA, A WRITE MAIN STORE OPERATION REQUIRES THAT THE DATA BE PLACED ON THE MIX BUS AFTER THE ADDRESS IS ESTABLISHED. THE SINGLE COMMAND MSGO IS ISSUED AS FOR THE READ MAIN STORE OPERATION. AT SOME TIME THEREAFTER, THE VALUE OF THE DATA IS ESTABLISHED BY THE COMBINATION OF A CHANGED VALUE IN FMIX OR THE REGISTER IT REFERENCES. THEN THE COMMAND MSRS IS ISSUED AND THE MEMORY CYCLE COMPLETES, PLACING THE NEW DATA ON THE MIX BUS INTO THE MEMORY AT THE ADDRESS PREVIOUSLY SPECIFIED.

IN ADDITION TO ALL THE REQUIREMENTS FOR READ MAIN STORE (SEE SECTION 5.4.3.1), ONE MUST ESTABLISH FMIX AND MIX NO LATER THAN THE T-STEP BEFORE MSRS IS ISSUED.

THE MEMORY CYCLE TIME IS VARIABLE, DEPENDING ON WHEN MSRS IS ISSUED. IN PARTICULAR, THE LEADING EDGE OF THE NEXT MSGO SHOULD NOT OCCUR LESS THAN 400 NANOSECONDS AFTER THE TRAILING EDGE OF MSRS OR LESS THAN 800 NANOSECONDS AFTER THE LEADING EDGE OF THE PREVIOUS MSGO.

5.4.3.3 READ-MODIFY-WRITE

MAIN STORE MAY BE OPERATED IN A READ-MODIFY-WRITE MODE IN ORDER TO MODIFY THE CONTENTS OF A MEMORY LOCATION BASED ON THE VALUE READ. THIS OPERATION IS DONE BY STARTING A WRITE MAIN STORE, WAITING FOR DATA AVAILABLE, EXTRACTING THE DATA WITH GATE MS, AND FINALLY ISSUING MSRS WHEN THE MODIFIED DATA IS READY ON THE MIX BUS TO BE WRITTEN BACK INTO MEMORY. THUS THE ONLY DIFFERENCE BETWEEN A NORMAL WRITE MAIN STORE AND READ-MODIFY-WRITE IS IN THE TIMING OF THE MSRS SIGNAL. THE TIMING RULES OF SECTIONS 5.4.3.1 AND 5.4.3.2 APPLY.

5.4.4 EXTERNAL OPERATIONS

THE COMMANDS FOR EXTERNAL OPERATIONS ARE TREATED IN THIS SECTION BECAUSE, FROM A PROGRAMMING POINT OF VIEW, THEY ARE QUITE SIMILAR TO THE COMMANDS FOR INITIATING VARIOUS MEMORY REFERENCE OPERATIONS.

TWO NANOPRIMITIVES ARE AVAILABLE FOR INITIATING EXTERNAL OPERATIONS. THESE ARE RIO ("READ IO") AND XIO ("TRANSMIT IO").

EXECUTION OF THE COMMAND RIO REQUIRES THAT THE PORT BE PREVIOUSLY SELECTED BY SETTING THE APPROPRIATE VALUE IN KA. THE RIO COMMAND CAUSES THE FOLLOWING ACTION:

1. THE PORT-RIO PULSE IS SENT IMMEDIATELY (LEADING EDGE).
2. THE PORT REGISTER (IN EXTERNAL STORE) IS CLEARED.
3. THE 18 INCOMING DATA LINES ARE GATED INTO THE PORT REGISTER.

EXECUTION OF THE COMMAND XIO ALSO REQUIRES THAT THE PORT BE PREVIOUSLY SELECTED IN KA. IF REQUIRED FOR THE EXTERNAL OPERATION, THE DATA ON THE G-BUS AND ON THE INPUT TO THE "PHANTOM" REGISTER (FIPH) MUST ALSO BE ESTABLISHED. THE XIO COMMAND CAUSES AN XIO PULSE TO BE SENT TO THE SELECTED PORT.

ADDITIONAL INFORMATION ON THE EXTERNAL INTERFACE IS GIVEN IN SECTION 4.6 AND SECTION 8.

5.5 DATA TRANSFER FUNCTIONS

5.5.1 T TRANSFERS

ALL T-TRANSFERS ARE INVOLVED WITH ESTABLISHING THE PRIMARY CONTROL STATE OF THE MACHINE. THUS SUCH TRANSFERS TAKE PLACE ON THE TRAILING EDGE T-CLOCK. THE COMMANDS CAUSING T-TRANSFERS ARE:

LOAD NPC - CAUSES A TRANSFER OF 10 BITS INTO THE NANOPROGRAM COUNTER (NPC) FROM ONE OF THREE SOURCES AS SPECIFIED:
 (KN) - 10 BITS FROM THE KN FIELD OF THE ACTIVE K-VECTOR.
 (SEQUENCE) - CURRENT VALUE IN NPC PLUS ONE MODULO 1024.
 (CS) - LOW ORDER 3 BITS OF FIDX PLUS HIGH ORDER SEVEN BITS OF COD BUS IN THE FOLLOWING ORDER:

FIDX	COD BUS
-----	-----
/ X X X A B C /	E F G H I J D X X X X X X X X X X /
-----	-----

-----	LOAD NPC (CS) ALSO CAUSES THE
NPC / A B C D E F G H I J /	LOW ORDER 11 BITS FROM THE COD
-----	BUS TO BE SAVED IN A DEDICATED
	REGISTER FOR FUTURE USE.

LOAD R31 - CLEARS THE HIGH ORDER SEVEN BITS OF R31 AND CAUSES A TRANSFER, INTO THE A AND B FIELDS OF R31, OF THE LOW ORDER 11 BITS FROM THE COD BUS THAT WAS SAVED BY THE MOST RECENT LOAD NPC (CS).

GATE NS UNCONDITIONALLY -CAUSES A TRANSFER OF THE 360 BITS LAST READ FROM NANOSTORE TO BE GATED INTO THE CONTROL MATRIX WITH EXECUTION TO BEGIN IN T1 OF THE WORD GATED, UNLESS T1 IS SKIPPED AS A RESULT OF SKIP ISSUED CONCURRENTLY WITH GATE NS UNCONDITIONALLY.

GATE NS - SAME AS GATE NS UNCONDITIONALLY PROVIDED THE TEST SPECIFICATION IN THE CURRENT T-STEP IS SATISFIED. OTHERWISE, THE COMMAND IS IGNORED. SEE SECTION 5.7.

SKIP - CAUSES ALL OF THE T-VECTOR COMMANDS IN THE NEXT T-STEP TO BE IGNORED IF THE TEST SECIFICATION IN THE CURRENT T-STEP IS SATISFIED. OTHERWISE THE COMMANDS WILL BE EXECUTED NORMALLY. GATE NS AND SKIP ARE MUTUALLY EXCLUSIVE IN A T-STEP.

5.5.2 F TRANSFERS

SIX-BIT TRANSFERS INTO AND OUT OF F-STORE (INCLUDING F INCREMENT, F DECREMENT AND ALUF OPERATIONS; SEE SECTION 5.6.6 AND 5.6.7) ARE CONTROLLED BY THREE GROUPS OF FIELDS WITHIN THE ACTIVE T-VECTOR:

GROUP 0: [FSEL0] [AUX0] [IN0] [OUT0]
 5 3 1 1 (BITS)

GROUP 1: [FSEL1] [AUX1] [IN1] [OUT1]
 5 3 1 1 (BITS)

GROUP 2: [FSEL2] [AUX2] [IN2]
 [AUX3] [OUT2]
 5 3 1 1 (BITS)

SINCE THE IN AND OUT CONTROLS SERVE AS NANOPRIMITIVE COMMANDS, SIX (AT MOST) SUCH TRANSFERS MAY OCCUR IN A SINGLE T-STEP. THE SIX COMMANDS ARE DEFINED AS FOLLOWS:

IN0: AUX [AUX0, WITH ITS SOURCE-AUX ENCODING] ---> F [FSEL0]
 OUT0: F [FSEL0] ---> AUX [AUX0, WITH ITS DESTINATION-AUX ENCODING]
 IN1: AUX [AUX1, WITH ITS SOURCE-AUX ENCODING] ---> F [FSEL1]
 OUT1: F [FSEL1] ---> AUX [AUX1, WITH ITS DESTINATION-AUX ENCODING]
 IN2: AUX [AUX2, WITH ITS SOURCE-AUX ENCODING] ---> -F [FSEL2]
 OUT2: F [FSEL2] ---> AUX [AUX3, WITH ITS DESTINATION-AUX ENCODING]

NORMALLY, ONLY TWO OR THREE OF THE F TRANSFERS ARE COMMANDED IN A GIVEN T-STEP SINCE THE F SELECTION MUST BE COMMON FOR EACH IN/OUT TRANSFER. THE AUX ENCODINGS ARE SHOWN SCHEMATICALLY IN THE DIAGRAM BELOW.

ENCODING OF F TRANSFERS

FIGURE 5.5.2A

GROUP 0				GROUP 1				GROUP 2						
-----				-----				-----						
I	SOURCE	[AUX0]	I	I	SOURCE	[AUX1]	I	I	SOURCE	[AUX2]	I			
I	0	000	A	I	I	0	000	A	I	I	0	000	A	I
I	1	001	B	I	I	1	001	B	I	I	1	001	B	I
I	2	010	****	I	I	2	010	C	I	I	2	010	KX	I
I	3	011	KA	I	I	3	011	KA	I	I	3	011	KA	I
I	4	100	KB	I	I	4	100	KT	I	I	4	100	KB	I
I	5	101	G(GSPEC)	I	I	5	101	G(GSPEC)	I	I	5	101	G(GSPEC)	I
I	6	110	ALUF	I	I	6	110	INCF1	I	I	6	110	INCF2	I
I	7	111	IO ID	I	I	7	111	DECF1	I	I	7	111	DECF2	I
-----				-----				-----						
	I				I				I					
	IN0	I			IN1	I			IN2	I				
	V				V				V					
-----				-----				-----						
	I	F	(FSEL0)	I		I	F	(FSEL1)	I		I	F	(FSEL2)	I
-----				-----				-----						
	I				I				I					
	OUT0	I			OUT1	I			OUT2	I				
	V				V				V					
-----				-----				-----						
I	DEST	[AUX0]	I	I	DEST	[AUX1]	I	I	DEST	[AUX3]	I			
I	0	000	A	I	I	0	000	A	I	I	0	000	A	I
I	1	001	B	I	I	1	001	B	I	I	1	001	B	I
I	2	010	C	I	I	2	010	C	I	I	2	010	C	I
I	3	011	KA	I	I	3	011	KA	I	I	3	011	KA	I
I	4	100	KB	I	I	4	100	KB	I	I	4	100	KB	I
I	5	101	KSHC	I	I	5	101	KX	I	I	5	101	KSHC	I
I	6	110	KALC	I	I	6	110	KALC	I	I	6	110	KT	I
I	7	111	KS	I	I	7	111	KSHA	I	I	7	111	KSHA	I
-----				-----				-----						

NOTES: G(GSPEC) IS DEFINED BELOW IN THIS SECTION. TRANSFERS INVOLVING THIS AUX SOURCE MUST APPEAR IN STRETCHED T-STEPS.
 INCF1, DECF1, INCF2 AND DECF2 ARE THE RESULTS OF THE F INCREMENT AND F DECREMENT. THEY ARE DESCRIBED IN 5.6.6.
 ALUF IS DESCRIBED IN 5.6.6. IO ID IS DESCRIBED IN 4.6.

G SPECIFIER

LISTED AS A SOURCE IN F TRANSFERS IS THE QUANTITY G(GSPEC). RATHER THAN BEING A SINGLE SOURCE FIELD, G(GSPEC) INDICATES ONE OF 16 POSSIBLE SOURCES ON THE "G-BUS", SELECTED BY THE VALUE OF A SINGLE FOUR-BIT FIELD (GSPEC) IN THE ACTIVE T-VECTOR. THE NORMAL USE OF GSPEC IS TO INDIRECTLY SPECIFY A SIX BIT VALUE BY REFERENCING ONE OF THE G REGISTERS (THE LAST 12 F REGISTERS). HOWEVER THE COMPLETE USE OF GSPEC INCLUDES THE FOLLOWING SIX DISJOINT NANOPRIMITIVE OPERATIONS. THE G-SPECIFIER (GSPEC) IS USED:

1. TO SPECIFY A SOURCE FOR CERTAIN 6-BIT TRANSFERS;
2. TO SPECIFY ALUF INPUT SELECTION (5.6.7);
3. TO SPECIFY THE OPERAND IN ALL INC MPC OPERATIONS (5.6.5);
4. AS PART OF THE QM-1 EXTERNAL INTERFACE (4.6).
5. IN THE GENERATE/CLEAR INTERRUPT FACILITY (4.5.2.4).
6. TO SPECIFY THE DESTINATION OF INDEX ALU OPERATIONS (5.6.4).

HENCE, UNLESS A BIT PATTERN IN THE GSPEC CAN BE SHARED TO ADVANTAGE, THESE FIVE TYPES OF OPERATIONS MUST BE CONSIDERED MUTUALLY EXCLUSIVE IN A T-STEP.

WHEN INVOLVED IN OPERATIONS (1), (2), (4), (5) OR (6) ABOVE, THE GSPEC SELECTS ONE OF THE 12 G'S OR ONE OF FOUR "PSEUDO-G'S", AS FOLLOWS:

VALUE OF GSPEC	G REGISTER SELECTED	VALUE OF GSPEC	PSUEDO G SELECTED	COMMAND SPECIFICATION
0	F20=G0	12	IO ID	"G IO ID"
1	F21=G1	13	B	"G B"
2	F22=G2	14	KS	"G KS"
3	F23=G3	15	KX	"G KX"
4	F24=G4			
5	F25=G5			
6	F26=G6			
7	F27=G7			
8	F28=G8			
9	F29=G9			
10	F30=G10			
11	F31=G11			

THE TERM, G(GSPEC) HAS
BEEN USED TO REFER TO
ANY OF THESE 16 SELECTIONS.

WHEN GSPEC IS USED IN (3) THE LOW ORDER TWO BITS ARE USED DIRECTLY TO SPECIFY THE INC MPC OPERAND (SEE SECTION 5.6.5).

5.5.3 R TRANSFERS

ALL OF THE TRANSFERS TO AND FROM LOCAL STORE AND EXTERNAL STORE ARE CLASSIFIED AS R TRANSFERS SINCE THEY OCCUR ON THE TRAILING EDGE R-CLOCK. ACTUALLY SUCH TRANSFERS ARE DIVIDED INTO TWO CLASSES; INPUT AND OUTPUT, AS VIEWED FROM THE STORE INVOLVED. ALL OUTPUT TRANSFERS ARE ENABLED AS SOON AS THE BUS CONTROL IS ESTABLISHED; INPUT TRANSFERS REQUIRE AN EXPLICIT COMMAND TO ENABLE THE GATING OF THE DATA.

THUS OUTPUT FROM LOCAL STORE OR EXTERNAL STORE IS ENABLED ON ALL OF THE FOLLOWING BUSES AS SOON AS THE CORRESPONDING F REGISTER SELECTS ONE OF THE APPROPRIATE REGISTERS:

BUS	CONTROL	DATA USED WHEN:
MIX	FMIX	MSGO OR MSRS
CIA	FCIA	READ CS OR WRITE CS
CID	FCID	WRITE CS
EID	FEID	LOAD ES
AIL	FAIL	ALU ALWAYS OPERATING
AIR	FAIR	" " "
SID	FSID	SHIFTER ALWAYS OPERATING
INC	FMPC ***	MPC AND INDEX FACILITY ALWAYS OPERATING
EOD	FEOD	GATE ES

INPUT TO LOCAL STORE OR EXTERNAL STORE REQUIRES NOT ONLY A BUS CONTROL BUT ALSO A COMMAND TO CAUSE THE ACTUAL GATING. FOR THE SIX INPUT BUSES TO LOCAL STORE AND THE INPUT TO EXTERNAL STORE, THESE ARE:

BUS	CONTROL	DATA TRANSFERRED WHEN
MOD	FMOD	GATE MS
COD	FCOD	GATE CS
EOD	FEOD	GATE ES
ROD	FROD	GATE ALU
SOD	FSOD	GATE SH
INC	FMPC ***	INC MPC OR INDEX
EID	FEID	LOAD ES

NOTE: ***-INDEX USES ANOTHER SELECTION MECHANISM, SEE SECTION 5.6.4.

THE INPUT TO ALL TRANSFERS IS DETERMINED BY THE STATE OF THE CONTROLS AS THEY EXISTED FOLLOWING THE R-CLOCK IN THE T-STEP IN WHICH THE TRANSFER IS ISSUED. THE TRANSFER IS COMPLETED FOLLOWING THE R CLOCK OF THE NEXT T-STEP.

5.6 DATA MANIPULATION FUNCTIONS

5.6.1 GENERAL

A VARIETY OF DATA MANIPULATION FUNCTIONS ARE PROVIDED IN THE QM-1. SECTION 5.6 TREATS EACH OF THE MAJOR COMPONENTS AVAILABLE FOR MODIFYING DATA IN EITHER THE 18-BIT OR 6-BIT DOMAINS. THESE COMPONENTS ARE:

ARITHMETIC-LOGIC UNIT (ALU)

SHIFTER AND SHIFTER EXTENSION

INDEX ALU

MPC FACILITY

F REGISTER INCREMENT AND DECREMENT FACILITY

ALUF

RMI

THE FIRST FOUR OPERATE ON DATA IN THE 18-BIT PORTION OF THE MACHINE. THE NEXT TWO ARE FOR MANIPULATION SIX-BIT DATA. THE LAST COMPONENT PROVIDES OPTIONAL MANIPULATIVE CAPABILITY APPLIED TO DATA ARRIVING FROM MAIN STORE.

IN ADDITION, R31 PROVIDES AN INTERFACE BY WHICH 18-BIT DATA MAY BE MANIPULATED IN SIX-BIT SUB-FIELDS. FINALLY, AN "AUTO OR" CAPABILITY EXISTS FOR LOGICALLY COMBINING TWO VALUES IN EITHER DOMAIN SIMULTANEOUS WITH A TRANSFER.

ALL OF THE DATA MANIPULATING COMPONENTS DESCRIBED IN THE FOLLOWING SECTIONS MAY BE OPERATED IN PARALLEL. THEY COMBINE TO PROVIDE AN EXTREMELY POWERFUL ARITHMETIC AND LOGICAL CAPABILITY.

5.6.2 OPERATION OF THE ARITHMETIC-LOGIC UNIT

THE ARITHMETIC-LOGIC UNIT (ALU) OUTPUT IS DETERMINED AS A FUNCTION OF:

- 18 BIT INPUT ON THE ALU INPUT LEFT (AIL) BUS.
- 18 BIT INPUT ON THE ALU INPUT RIGHT (AIR) BUS.
- A SINGLE BIT, THE CARRY-IN-HOLD (CIH) FLIPFLOP.
- A 6 BIT VALUE IN THE K REGISTER (KALC) USED TO CONTROL THE ALU.
- A SINGLE BIT IN THE F REGISTER (FIDX) USED TO SPECIFY 16 OR 18 BIT MODE.

THE BASIC FUNCTION IS DETERMINED BY THE LOW ORDER 4 BITS OF KALC AS FOLLOWS:

KALC	1	BIT	1	BIT	1	BIT	1	BIT	1	BIT	1
LAYOUT	1	5	1	4	1	3	1	2	1	1	0

DECIMAL LOGIC											
CONTROL CONTROL <-----BASIC ALU FUNCTION----->											
CODE CODE											

ALU OUTPUT TRUTH TABLE

THE INDIVIDUAL BIT OUTPUTS ARE DETERMINED BY THE TWO INPUT BIT VALUES AND THE 4 LOW ORDER BITS OF KALC AS INDICATED IN THE TRUTH TABLE ON THE RIGHT. FOR EXAMPLE, THE EXCLUSIVE OR FUNCTION RESULTS WHEN KALC IS 011001. THIS CAN BE SEEN BY INSERTING THE CONTROL BITS INTO THE TRUTH TABLE AS SPECIFIED. THUS ALL OF THE 16 BOOLEAN FUNCTIONS ARE AVAILABLE AND THEIR ENCODING CAN BE DEDUCED FROM THE TRUTH TABLE GIVEN.

		RIGHT INPUT	
		0	1
L		-----	
E		-----	
F		I NOT I	NOT I
T	0	I BIT I	BIT I
		I 3	I 2
I		-----	
N		I	I
P	1	I BIT I	BIT I
U		I 0	I 1
T		-----	

IN ALL OF THE CASES ABOVE, THE LOGIC CONTROL CODE IS SET TO 1 TO DISABLE THE INTERNAL CARRY INTO EACH BIT. WHEN THE LOGIC CONTROL CODE IS 0, THE OPERATION ON EACH BIT IS MODIFIED BY THE CARRY INTO THAT BIT. THE CARRY INTO THE BIT IS EXCLUSIVE OR'ED WITH THE RESULT ABOVE TO PRODUCE THE MODIFIED RESULT FOR THE BIT. THIS IS SHOWN IN THE MODIFIED TRUTH TABLE BELOW.

HERE, CIN IS DEFINED AS THE CARRY INTO THE BIT IN QUESTION. FOR THE LOW ORDER BIT, CIN IS DEFINED AS THE VALUE OF THE CARRY IN HOLD (CIN) FLIPFLOP. FOR EACH SUBSEQUENT BIT, CIN IS THE SAME AS THE CARRY OUTPUT FROM THE PRECEDING BIT.

THE CARRY OUTPUT FROM ANY BIT DEPENDS ON THE INPUTS TO THE BIT (INCLUDING CIN) AND ON THE FUNCTION SPECIFIED AS BEFORE. THIS RELATIONSHIP IS GIVEN IN THE FINAL TRUTH TABLE BELOW.

ALU CARRY OUTPUT TRUTH TABLE

RIGHT INPUT					
		0		1	
L		-----			
E		-----			
F		I NOT BIT 3	I NOT BIT 2	I	
T 0	I	AND	I AND	I	
	I	CIN	I CIN	I	
I		-----			
N		I NOT BIT 0	I NOT BIT 1	I	
P 1	I	OR	I OR	I	
U	I	CIN	I CIN	I	
T		-----			

NORMAL OPERATION OF THE ALU IS IN 18 BIT MODE. THIS RESULTS WHEN THE 16 BIT MODE CONTROL (HIGH ORDER BIT IN FIDX) IS RESET. WHEN IN 18 BIT MODE, ALL 18 BITS ARE ACTIVE ON THE INPUTS TO THE ALU. HOWEVER, WHEN 16 BIT MODE CONTROL IS SET, BIT 15 ON EACH OF THE INPUTS (AIL AND AIR) IS REPLICATED TO FORM THE INPUTS TO BIT 16 AND BIT 17. THIS APPLIES TO ALL ALU OPERATIONS EXCEPT "PASS LEFT" (011111) WHICH PROVIDES A MEANS OF TRANSMITTING ALL 18 BITS THROUGH THE ALU EVEN WHEN IN 16 BIT MODE. OTHERWISE, THE ALU OPERATES IDENTICALLY IN EITHER 16 OR 18 BIT MODE.

ALU ARITHMETIC OUTPUT TRUTH TABLE

		RIGHT INPUT			
L		0		1	
E		-----			
F		I	NOT BIT 3	I	NOT BIT 2
T 0		I	XOR	I	XOR
		I	CIN	I	CIN
I		-----			
N		I	BIT 0	I	BIT 1
P 1		I	XOR	I	XOR
U		I	CIN	I	CIN
T		-----			

THE CARRY OUTPUT IS VALID INDEPENDENT OF THE SETTING OF THE LOGIC CONTROL CODE. THIS IS A CONSEQUENCE OF THE FACT THAT LOGICAL MODE ONLY INHIBITS THE EFFECT OF CARRY ON THE OUTPUT BIT AND CAUSES NO CHANGE IN CARRY PROPAGATION WITHIN THE ALU.

THIS DETAILED COVERAGE OF THE ALU HAS BEEN INCLUDED SO THAT ANY QUESTION ABOUT THE ALU OPERATION CAN BE ANSWERED. FOR NORMAL USE, THE TABLE ON THE NEXT PAGE SUFFICES FOR ENCODING ALL ALU OPERATIONS.

ALU CONTROL IS TAKEN FROM THE KALC FIELD IN THE ACTIVE K-VECTOR. THE HIGH-ORDER BIT OF KALC IS THE DECIMAL CONTROL; WHEN IT IS ON, DECIMAL CORRECTION WORDS ARE FORCED ONTO THE SOD BUS TO REFLECT CURRENT ALU ACTIVITY, AS PER SECTION 4.2.3.1. NOTE: WHEN THE DECIMAL CONTROL IS ON, KSHC MUST SPECIFY SINGLE SHIFTS SINCE THE DECIMAL CONTROL ALSO FORCES THE SHIFTER TO OPERATE IN SINGLE MODE INDEPENDENT OF KSHC (SEE SECTION 5.6.3).

THE REMAINING FIVE BITS OF KALC CONTROL ALU FUNCTION AS FOLLOWS:
(L,R=LEFT,RIGHT; 2'S COMPLEMENT CONVENTION ASSUMED)

BIT 3 THRU BIT 0 OF KALC	LOGICAL FUNCTIONS KALC BIT 4 = 1	ARITHMETIC FUNCTIONS KALC BIT 4 = 0		CARRY-OUT BIT IS DEFINED AS CARRY-IN BIT
		CARRY IN HOLD = 0	CARRY IN HOLD = 1	
0000	NOT L	$L - 1$	L	OR L
0001	NOT (L AND R)	$\langle L \text{ AND } R \rangle - 1$	L AND R	OR $\langle L \text{ AND } R \rangle$
0010	NOT L OR R	$\langle L \text{ AND NOT } R \rangle - 1$	L AND NOT R	OR $\langle L \text{ AND NOT } R \rangle$
0011	ALL ONES	ALL ONES	ALL ZEROS	OR ZERO
0100	NOT (L OR R)	$\langle L \text{ OR NOT } R \rangle + L$	$\langle L \text{ OR NOT } R \rangle + L + 1$	
0101	NOT R	$\langle L \text{ OR NOT } R \rangle +$ $\langle L \text{ AND } R \rangle$	$\langle L \text{ OR NOT } R \rangle +$ $\langle L \text{ AND } R \rangle + 1$	
0110	NOT (L XOR R)	$L - R - 1$	$L - R$	
0111	L OR NOT R	$\langle L \text{ OR NOT } R \rangle$	$\langle L \text{ OR NOT } R \rangle + 1$	AND $\langle L \text{ OR NOT } R \rangle$
1000	NOT L AND R	$\langle L \text{ OR } R \rangle + L$	$\langle L \text{ OR } R \rangle + L + 1$	
1001	L XOR R	$L + R$	$L + R + 1$	
1010	R	$\langle L \text{ OR } R \rangle +$ $\langle L \text{ AND NOT } R \rangle$	$\langle L \text{ OR } R \rangle +$ $\langle L \text{ AND NOT } R \rangle + 1$	
1011	L OR R	L OR R	$\langle L \text{ OR } R \rangle + 1$	AND $\langle L \text{ OR } R \rangle$
1100	ALL ZEROS	$L + L$	$L + L + 1$	
1101	L AND NOT R	$L + \langle L \text{ AND } R \rangle$	$L + \langle L \text{ AND } R \rangle + 1$	
1110	L AND R	$L + \langle L \text{ AND NOT } R \rangle$	$L + \langle L \text{ AND NOT } R \rangle + 1$	
1111	L [PASS LEFT]	L	$L + 1$	AND L

THE OUTPUT OF THE ALU CONSISTS OF THE FOLLOWING:

AN 18 BIT OUTPUT DETERMINED BY THE ABOVE RULES.

A CARRY-OUT CONDITION (CARRY OUT OF THE HIGH ORDER BIT).

AN OVERFLOW CONDITION DETERMINED BY THE EXCLUSIVE OR OF THE CARRY
OUT VALUES FROM THE TWO HIGH ORDER BITS.

THE 18-BIT OUTPUT VALUE IS USED AS THE 18-BIT INPUT TO THE SHIFTER EXTENSION
DISCUSSED IN SECTION 5.6.3. THIS VALUE MAY BE PASSED DIRECTLY TO THE ADD BUS
OR IT MAY BE SHIFTED BEFORE TRANSFER.

THE CARRY-OUT CONDITION MAY BE TRANSFERRED TO THE CARRY-OUT-HOLD FOR TESTING
AND IT MAY BE ALSO TRANSFERRED TO THE CARRY-IN-HOLD FOR FUTURE USE. SECTION
4.2.3.4 COVERS CARRY CONTROL. TESTING OF THE CARRY OUT AND OVERFLOW CONDITIONS
IS DISCUSSED IN SECTION 5.7.

IN GENERAL, OPERATION OF THE ALU REQUIRES 2 T-PERIODS; EITHER ONE STRETCHED
T-STEP OR TWO UNSTRETCHED T-STEPS. DETAILED TIMING OF BOTH THE ALU AND THE
SHIFTER IS COVERED IN SECTION 5.6.3.

5.6.3 OPERATION OF THE SHIFTER AND SHIFTER EXTENSION

THE SHIFTER AND SHIFTER EXTENSION FORM A COMBINED SHIFT MATRIX UNIT OPERATING ON A TOTAL OF 36 BITS ENTERING FROM THE ALU (HIGH ORDER 18 BITS) AND FROM THE SID BUS. THIS UNIT MAY OPERATE IN EITHER SINGLE MODE (PASSING THE ALU OUTPUT ON TO THE ADD BUS) OR DOUBLE MODE (SHIFTING THE ALU OUTPUT BEFORE IT REACHES THE ADD BUS). THE 36 BIT OUTPUT TO THE ADD AND SOD BUSES IS DETERMINED AS A FUNCTION OF:

18 BIT INPUT FROM THE ALU.

18 BIT INPUT FROM THE SID BUS.

A 6 BIT VALUE IN THE K REGISTER (KSHA) USED TO SPECIFY THE SHIFT AMOUNT.

A 6 BIT VALUE IN THE K REGISTER (KSHC) USED TO CONTROL THE SHIFTER.

A SINGLE BIT, THE CARRY-OUT-HOLD (COH) FLIPFLOP.

THE HIGH ORDER BIT (DECIMAL CONTROL CODE) OF THE ALU CONTROL, KALC.

KSHA SPECIFIES SHIFT AMOUNT (NUMBER OF POSITIONS), AND IS INTERPRETED MODULO X, WHERE X IS APPROPRIATE FOR THE TYPE OF SHIFT SPECIFIED IN KSHC.

KSHC IS INTERPRETED AS FOLLOWS:

KSHC LAYOUT	I	BIT	I	BIT	I	BIT	I	BIT	I	BIT	I	BIT	I
	I	5	I	4	I	3	I	2	I	1	I	0	I

	LEFT		RIGHT		TYPE				MODE		DIRECTION		
	CONTROL		CONTROL		00 CIRCULAR				0 SINGLE		0 LEFT		
	SWITCH		SWITCH		01 LOGICAL				1 DOUBLE		1 RIGHT		
					10 ARITHMETIC								
	(SECTION 4.2.3.4)				11 UNDEFINED								

TIMING OF ALU AND SHIFTER

ASSUME THAT THE LAST TRAILING-EDGE OPERATION THAT CHANGES ANY INPUT RELEVANT TO AN ALU/SHIFT PROCESS IS COMMANDED IN T-PERIOD T_N . THEN THE EARLIEST T-PERIOD (NOT T-STEP) IN WHICH THE DATA OUTPUTS OF THAT PROCESS ARE AVAILABLE AS DEFINED FOR GATING OUT BY A TRAILING-EDGE NANOPRIMITIVE IS GIVEN BY T_N+X , WHERE X DEPENDS ON THE PROCESS, AS FOLLOWS (TEST OUTPUTS ARE ALSO AVAILABLE IN T_N+X):

PASS DATA THROUGH ALU (KALC = "PASS LEFT"), THEN BYPASS THE SHIFTER EXTENSION TO ADD (KSHC SPECIFIES SINGLE SHIFT) ----- $X = 2$.

PASS DATA THROUGH ALU ("PASS LEFT"), THEN DOUBLE SHIFT ----- $X = 2$.

ALU OPERATION (NOT "PASS LEFT"), THEN BYPASS THE SHIFTER EXTENSION ----- $X = 2$.

ANY SINGLE SHIFT, OR SID TO SOD WITH NO SHIFT (KSHA = 0) ----- $X = 2$.

ALU OPERATION (NOT "PASS LEFT"), THEN DOUBLE SHIFT (SPECIAL CASE: SID AND SHIFT CONTROL INPUTS NOT NEEDED UNTIL T_N+1) ----- $X = 3$.

THE ALU AND SHIFTING UNITS ARE NOT PIPELINE DEVICES; INPUTS TO A PROCESS MUST BE HELD STABLE FOR THE DURATION OF THAT PROCESS. (A NEW PROCESS IS DEFINED TO BEGIN WHEN ANY INPUT CHANGES; AT THAT TIME, THE RESULT OF THE PREVIOUS PROCESS IS CONSIDERED INVALID.) IF THIS RULE IS NOT FOLLOWED, OUTPUTS ARE UNDEFINED.


5.6.4 OPERATION OF THE INDEX ALU

A SPECIAL DATA MANIPULATION UNIT, THE INDEX ALU, IS PROVIDED IN ORDER TO FACILITATE RAPID INDEXING AND LOGICAL OPERATIONS. IN PARTICULAR, THIS FEATURE OF THE MACHINE IS USEFUL IN COMPUTING ADDRESSES AND IN MASKING OPERATIONS, ALTHOUGH IT IS NOT RESTRICTED TO THESE TASKS IN ANY SENSE. OPERATION OF THE INDEX ALU FUNCTION CAN BE ACCOMPLISHED IN A SINGLE STRETCHED T STEP, OR IN 2 ADJACENT UNSTRETCHED T-STEPS. THE RESULT IS GATED INTO LOCAL STORE WITH THE NANOPRIMITIVE "INDEX" (SEE 5.2.2).

THE IMPLEMENTATION OF THE FUNCTION REQUIRES THE DEDICATED USE OF THE THIRD F TRANSFER MECHANISM. NAMELY THE FOLLOWING CONTROL MATRIX FIELDS ARE USED IN ANY T STEP IN WHICH THE FUNCTION IS INVOKED.

AUX2	SELECTS LOCAL STORE SOURCE REGISTER
FSEL2	SELECTS ARITHMETIC OR LOGICAL FUNCTION
AUX3	SELECTS INDEX REGISTER
GSPEC	SELECTS LOCAL STORE DESTINATION REGISTER

EXCEPT FOR THOSE FIELDS, THE FUNCTION OPERATES INDEPENDENTLY AND CONCURRENT WITH ALL OTHER FUNCTIONS OF THE QM-1. THE ILLUSTRATION BELOW SHOWS THE DATA PATHS THAT ARE INVOLVED. NOTE THAT A LOCAL STORE REGISTER ALWAYS RECEIVES THE RESULT OF THE OPERATION.



THE THREE REGISTER SELECTION CONTROLS USED IN INDEX ALU OPERATION ARE ENCODED AS FOLLOWS:

AUX2 - PROVIDES AN INDIRECT REFERENCE TO A LOCAL STORE REGISTER WHICH WILL BE USED AS THE LEFT INPUT TO THE INDEX ALU. THE FIELD IS DECODED BY THE SAME PHYSICAL HARDWARE USED TO DECODE AUX2 IN ITS PRIMARY CAPACITY AS AN INPUT SELECT FOR AN F REGISTER TRANSFER. THEREFORE THE VALUES ASSOCIATED WITH EACH VALUE OF AUX2 ARE THE SAME AS WHEN IT IS USED TO SELECT AN AUX TO F TRANSFER.	AUX2	FIELD SELECTED
	0	A
	1	B
	2	KX
	3	KA
	4	KB
	5	G(GSPEC)
	6	NOT USED
	7	NOT USED

THE SELECTED FIELD, A OR B OR KX ETC. PROVIDES THE LOCAL STORE REGISTER SOURCE SELECTION. IF USED, GSPEC MUST BE REPEATED IN THE PREVIOUS T-STEP.

AUX3 - AN INDIRECT REFERENCE TO ONE OF 16 POSSIBLE SOURCES OF AN 18 BIT INDEX OPERAND WHICH WILL BE USED AS THE RIGHT INPUT TO THE INDEX ALU. TWELVE OF THESE ARE EXTERNAL REGISTERS; TWO ARE SOURCES OF ALL ONES; THE REMAINING TWO ARE THE MEMORY BUSES MOD AND COD RESPECTIVELY.	AUX3 SELECTS	CONTENTS SELECT "X" OPERAND
	0 A	0 XX0000 E8
	1 B	1 XX0001 E9
	2 KT	- - - - -
	3 KB	11 XX1011 E19
	4 F28=G8	12 XX1100 ALL ONES
	5 F29=G9	13 XX1101 ALL ONES
	6 F30=G10	14 XX1110 MOD
	7 F31=G11	15 XX1111 COD

GSPEC - AN INDIRECT REFERENCE TO A LOCAL STORE REGISTER TO RECEIVE THE RESULT OF THE INDEX ALU OPERATION. THE FIELD IS DECODED BY THE SAME HARDWARE USED TO DECODE GSPEC WHEN IT IS USED TO SELECT AN INPUT FOR A F REGISTER TRANSFER. THUS THE FIELD SELECTED BY GSPEC IS THE SAME AS WHEN IT IS USED TO SELECT AN AUX TO F TRANSFER.	GSPEC	FIELD SELECTED
	0	F20=G0
	-	- - - - -
	11	F31=G11
	12	NOT USED
	13	B
	14	KS
	15	KX

THE CONTROLS PROVIDING FOR SELECTION OF A LOCAL STORE REGISTER SOURCE AND DESTINATION AND FOR SELECTION OF AN INDEX OPERAND FROM EITHER AN E REGISTER OR DIRECTLY FROM MAIN STORE OR CONTROL STORE HAVE BEEN DESCRIBED. THERE REMAINS THE MANNER IN WHICH THE SPECIFIC FUNCTION IS SELECTED BY THE FSEL2 FIELD.

FSEL2 - EITHER A DIRECT SELECTION OF ONE OF 13 ARITHMETIC AND LOGICAL FUNCTIONS, OR AN INDIRECT SPECIFICATION OF ONE OF 18 POSSIBLE 6 BIT FIELDS CONTAINING ONE OF THE STANDARD 48 FUNCTION CODES. ("L" REFERS TO A LOCAL STORE REGISTER; "X" REFERS TO AN EXTERNAL STORE REGISTER, MOD BUS, OR COD BUS.

FSEL2	FUNCTION SPECIFIED	FSEL2	FUNCTION SPECIFIED BY	OR FUNCTION SPECIFIED
00	L - 1 ----> L	12	A	
01	L + 1 ----> L	13	B	
02	L XOR X ----> L	14	KA	
03	ALL ONES --> L	15	KB	
04	ALL ZERO --> L	16	F16 - FMPC	
05	NOT X ----> L	17	F17 - FIDX	
06	L - X ----> L	18		L --> L
07	L AND X ----> L	19		L --> L
08	L OR X ----> L	20	F20 - G0	
09	L + X ----> L	--	--	
10	X ----> L	31	F31 - G11	
11	NOT L ----> L			

FUNCTION CODE		FUNCTION SPECIFICATION	
BIT	I	5 4	5 4
3 2 1 0	I	* 1 (*=DON'T CARE)	0 0
-----		-----	
0 0 0 0	I	NOT L	I L - 1
0 0 0 1	I	NOT (L AND X)	I (L AND X) - 1
0 0 1 0	I	NOT L OR X	I (L AND NOT X) - 1
0 0 1 1	I	ALL ONES	I - 1 (2'S COMPLEMENT)
0 1 0 0	I	NOT (L OR X)	I L + (L OR NOT X)
0 1 0 1	I	NOT X	I (L AND X) + (L OR NOT X)
0 1 1 0	I	NOT (L XOR X)	I L - X - 1
0 1 1 1	I	L OR NOT X	I L - X
1 0 0 0	I	NOT L AND X	I (L OR NOT X) + 1
1 0 0 1	I	L XOR X	I L + (L OR X) + 1
1 0 1 0	I	X	I L + X + 1
1 0 1 1	I	L OR X	I (L AND NOT X) + (L OR X) + 1
1 1 0 0	I	ALL ZERO	I (L OR X) + 1
1 1 0 1	I	L AND NOT X	I L + L + 1
1 1 1 0	I	L AND X	I (L AND X) + L + 1
1 1 1 1	I	L	I (L AND NOT X) + L + 1
			I L + 1

5.6.5 MPC OPERATIONS

FOUR LOCAL STORE REGISTERS (R24, R25, R26 AND R27) HAVE A SPECIAL INCREMENT CAPABILITY TO FACILITATE THEIR USE AS MICRO PROGRAM COUNTERS. THE REGISTER CURRENTLY DESIGNATED AS THE MPC IS DETERMINED BY THE CONTENTS OF FMPC (MOD 4).

THE CURRENT CONTENT OF THE DESIGNATED MPC IS CONTINUOUSLY BEING ADDED TO THE FOLLOWING FOUR VALUES:

+1

+2

B (LOW ORDER 6 BITS OF R31, SIGN EXTENDED, TWO'S COMPLEMENT ADDITION).

AB (LOW ORDER 11 BITS OF R31, SIGN EXTENDED, TWO'S COMPLEMENT ADDITION).

THE RESULTS OF THESE COMPUTATIONS MAY BE USED AS A CONTROL STORE ADDRESS FOR READ CS OR WRITE CS, OR THEY MAY BE GATED BACK INTO THE DESIGNATED MPC, USING THE COMMAND INC MPC. THE CS ADDR SELECT FIELD IS USED TO SELECT THE DESIRED VALUE FOR CONTROL STORE OPERATIONS. THE GSPEC FIELD IS USED TO SELECT THE DESIRED RESULT IN INC MPC OPERATIONS. THE ENCODINGS ARE AS FOLLOWS:

MPC RESULT	CS ADDR SELECT	GSPEC VALUE
MPC	010	(X=IGNORED)
MPC+1	011	XX00
MPC+2	100	XX01
MPC+B	101	XX10
MPC+AB	110	XX11

ALL MPC OPERATIONS INVOLVE 16 BIT ARITHMETIC; THE HIGH ORDER TWO BITS ARE SIGN EXTENDED FROM THE 15TH BIT. THUS INCREMENTING $2^{16} - 1$ BY 1 PRODUCES ALL ONES (-1 IN TWO'S COMPLEMENT NOTATION). FOR CONTROL STORE OPERATIONS, THE HIGH ORDER THREE BITS ARE IGNORED.

THE INC MPC NANOPRIMITIVE IS A TRAILING-EDGE OPERATION; THE INCREMENTED VALUE IS LOADED INTO THE APPROPRIATE LOCAL STORE MPC (THE ONE SELECTED BY FMPC) AT THE END OF THE T-STEP IN WHICH THE NANOPRIMITIVE IS EXECUTED. THE PROPAGATION TIME THROUGH THE ADDING CIRCUITS IS SUCH THAT IF FMPC, MPC, A OR AB IS CHANGED BY ANY TRAILING-EDGE COMMAND IN THE T-STEP PRIOR TO THE ONE CONTAINING THE INC MPC COMMAND, THE T-STEP CONTAINING THE INC MPC COMMAND MUST BE STRETCHED.

5.6.6 INCREMENT F AND DECREMENT F OPERATIONS

A FACILITY IS AVAILABLE FOR INCREMENTING OR DECREMENTING THE CONTENTS OF ANY F REGISTER. THIS IS ACCOMPLISHED BY SELECTING THE DESIRED F REGISTER FOR AN AUX TO F TRANSFER, SELECTING INCF OR DECF AS THE AUX SOURCE AND ENABLING THE TRANSFER WITH THE IN COMMAND. SINCE INCF AND DECF ARE AVAILABLE AS AUX SOURCES ONLY IN GROUP 1 AND GROUP 2, AT MOST TWO F-REGISTERS MAY BE LOADED WITH (EITHER) AN INCREMENTED OR DECREMENTED VALUE IN A T-STEP.

THE F REGISTER INCREMENT AND DECREMENT FACILITY OPERATES BY CONTINUOUSLY PERFORMING AN INCREMENT AND DECREMENT OPERATION ON THE CONTENTS OF THE F REGISTERS SELECTED BY THE FSEL1 AND FSEL2 FIELDS. FOUR VALUES ARE AVAILABLE AS OUTPUTS:

$$F(FSEL1) + 1 \quad F(FSEL1) - 1 \quad F(FSEL2) + 1 \quad F(FSEL2) - 1$$

THESE VALUES ARE SELECTED BY THE CORRESPONDING AUX1 AND AUX2 FIELDS. THE COMMANDS IN1 AND IN2 CAUSE THE SELECTED VALUE TO BE GATED BACK INTO THE F REGISTER AS A TRAILING-EDGE F TRANSFER. THE F REGISTER INCREMENT AND DECREMENT ARE UNSIGNED SIX-BIT OPERATIONS. FOR EXAMPLE, INCREMENTING 111111 PRODUCES 000000.

BECAUSE OF THE PROPAGATION TIME OF THE ADDING CIRCUITS, AN INCF OR DECF MUST BE PLACED IN A STRETCHED T-STEP IN ORDER TO BE COMPLETELY SELF-CONTAINED. HOWEVER, IT MAY BE SUCCESSFULLY PLACED IN AN UNSTRETCHED T-STEP PROVIDED THE FOLLOWING TWO CONDITIONS ARE MET:

- 1) THE VALUE IN THE F REGISTER IS NOT CHANGED BY AN F TRANSFER IN THE PRECEDING T-STEP.
- 2) THE SAME F REGISTER SELECTION (FSEL1 OR FSEL2) IS SPECIFIED IN THE PRECEDING T-STEP.

NOTE: INCREMENTING OR DECREMENTING FIPH SHOULD BE AVOIDED SINCE THIS CONNECTS AN ADDER'S OUTPUT DIRECTLY TO ONE OF ITS INPUTS AND RESULTS IN AN UNSTABLE CONDITION.

5.6.7 ALUF OPERATIONS

A SIX-BIT ARITHMETIC LOGICAL FACILITY, DESIGNATED "ALUF", IS AVAILABLE AS A QM-1 OPTION. THIS FACILITY PROVIDES FULL ARITHMETIC AND LOGICAL OPERATIONS, USING THE CONTENTS OF TWO F REGISTERS AS THE TWO OPERANDS, WITH THE RESULT PLACED IN ANY DESIRED F REGISTER.

THE ALUF OUTPUT IS TRANSFERRED INTO AN F-REGISTER BY THE NORMAL AUX TO F TRAILING-EDGE NANOPRIMITIVES THAT LOADS SOURCE AUX FIELDS INTO AN F-REGISTER; THESE CONTROLS ARE PRESENTED IN SECTION 5.5.2. THIS TRANSFER IS ONLY AVAILABLE IN GROUP 0. FSEL0 DETERMINES THE F REGISTER RECEIVING THE RESULT.

THE LEFT DATA INPUT IS THE CONTENTS OF THE F REGISTER SELECTED BY FSEL1. THIS IS DESIGNATED "F" IN THE TABLE BELOW. THE RIGHT DATA INPUT IS FROM THE G REGISTER SELECTED BY THE GSPEC FIELD. THIS OPERAND IS REFERENCED AS "G" IN THE TABLE BELOW. FINALLY, THE ARITHMETIC OR LOGICAL OPERATION TO BE PERFORMED ON THESE TWO OPERANDS IS DETERMINED BY THE CONTENTS OF FSEL2 AS ENCODED IN THE TABLE BELOW.

ALUF FUNCTION ENCODING NOTE: 2'S COMPLEMENT CONVENTION ASSUMED.

LOW 4 BITS OF FSEL2	ALUF OUTPUT FSEL2 HIGH BIT = 1	ALUF OUTPUT FSEL2 HIGH BIT = 0
0000	NOT F	F
0001	NOT (F AND G)	F AND G
0010	NOT F OR G	F AND NOT G
0011	ALL ONES *****	ALL ZEROS
0100	NOT (F OR G)	F PLUS (F OR NOT G) PLUS 1
0101	NOT G	F AND G PLUS (F OR NOT G) PLUS 1
0110	NOT (F XOR G)	F MINUS G
0111	F OR NOT G	(F OR NOT G) PLUS 1
1000	NOT F AND G	F PLUS (F OR G)
1001	F XOR G	F PLUS G
1010	G	F AND NOT G PLUS (F OR G)
1011	F OR G	F OR G
1100	ALL ZEROS	F PLUS F
1101	F AND NOT G	(F AND G) PLUS F
1110	F AND G	(F AND NOT G) PLUS F
1111	F	F

***** DEFAULT CONDITION WHEN ALUF OPTION IS NOT INSTALLED.

THE ALUF UNIT IS CONTINUOUSLY OPERATING, USING THE INPUTS AND FUNCTION AS SPECIFIED BY THE FSEL1, GSPEC AND FSEL2 FIELDS AT THE LEADING-EDGE OF EACH T-STEP. THE OUTPUT WILL BE GATED AS A TRAILING-EDGE AUX TO F TRANSFER ONLY WHEN AUX0 SPECIFIES ALUF AND THE IN0 CONTROL SPECIFIES AN AUX TO F TRANSFER IN GROUP 0. ANY T-STEP USING THE ALUF OUTPUT MUST BE STRETCHED IN ORDER FOR THE OUTPUT TO BE STABLE.

IF THE ALUF OPTION IS NOT INSTALLED, A SOURCE OF ALL ONES IS AVAILABLE BY SELECTING THE APPROPRIATE ALUF CONTROLS.

5.6.8 RMI OPERATIONS

A SPECIAL UNIT, DESIGNATED THE RMI UNIT, IS AVAILABLE AS A QM-1 OPTION, FOR MANIPULATING THE DATA READ FROM MAIN STORE ON THE WAY INTO LOCAL STORE OR EXTERNAL STORE. THIS UNIT PERMITS SELECTION AMONG THREE SETS OF DYNAMICALLY VARIABLE ROTATE, MASK AND INDEX OPERATIONS ON THE MAIN STORE DATA. IT IS IDEALLY SUITED FOR EMULATION OF VARIOUS MAIN STORE MACHINES HAVING A VARIETY OF MEMORY WIDTHS AND FORMATS.

A PRELIMINARY DESCRIPTION OF THE RMI UNIT IS GIVEN IN SECTION 4.2.6.4. FURTHER SPECIFICATIONS ARE AVAILABLE ON REQUEST.

5.7. CONDITIONAL FUNCTIONS

5.7.1. TEST FIELDS AND MASKS

THREE CLASSES OF TEST CONDITIONS ARE AVAILABLE FOR DECISION MAKING AT THE NANO-PROGRAM LEVEL. THESE ARE:

- LOCAL CONDITIONS - SIX CURRENT OUTPUT CONDITIONS AVAILABLE FROM THE ALU AND SHIFTER.
- GLOBAL CONDITIONS - SIX OUTPUT CONDITIONS FROM THE ALU AND SHIFTER THAT WERE PREVIOUSLY SAVED IN F REGISTER FIST.
- SPECIAL CONDITIONS - A SET OF SIX SPECIAL MACHINE STATUS CONDITIONS THAT AVAILABLE FOR TESTING.

WITHIN ANY CLASS, ONE OR MORE OF THE CONDITIONS MAY BE SIMULTANEOUSLY TESTED FOR BEING PRESENT OR ABSENT. AN INDIVIDUAL SIX-BIT MASK FIELD IN THE ACTIVE K-VECTOR IS PROVIDED FOR TESTING EACH CLASS. THE BITS SET IN THE MASK SELECT THE CONDITIONS TO BE TESTED.

KT MASKS THE LOCAL CONDITIONS.
 KS MASKS THE GLOBAL CONDITIONS (FIST);
 KX MASKS THE SPECIAL CONDITIONS.

THE ARITHMETIC TEST FIELDS AND MASKS (KS, KT, AND FIST) ARE FORMATTED AS SHOWN BELOW. (THE CONDITIONS ARE DEFINED IN SECTION 4.2.3.5.)

I SHIFTER	I ALU	I ALU	I ALU	I COMBINED	I SHIFTER
I HIGH BIT	I CARRY	I SIGN	I RESULT	I OVERFLOW	I LOW BIT
-----	-----	-----	-----	-----	-----
BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0

THE SPECIAL TEST FIELD AND MASK (KX) IS FORMATTED AS FOLLOWS:

I NOT USED	I PROGRAM	I R INDEX	I MS	I MS DATA	I F NOT
I RESERVED	I CHECK	I NOT ZERO	I BUSY	I INVALID	I ZERO
-----	-----	-----	-----	-----	-----
BIT 5	BIT 4	BIT 3	BIT 2	BIT 1	BIT 0

SEE SECTION: 4.5.2.2 5.6.4 4.2.6 5.6.6

HENCE AT MOST ONE CLASS OF CONDITIONS MAY BE TESTED IN A T-STEP. HOWEVER, IT IS POSSIBLE TO TAKE ACTION UPON COMBINATIONS OF CONDITIONS WITHIN A CLASS; FOR EXAMPLE, A USEFUL TEST INVOLVES SPECIFYING BOTH THE S AND R BITS IN KT, THUS TESTING FOR A FULL 18-BIT OR 16-BIT ARITHMETIC RESULT.

5.7.2 TESTS

THE EXECUTION OF AN ACTION (SKIP OR GATE NS) AS THE RESULT OF A TEST IS CONSIDERED A TRAILING-EDGE EVENT. HOWEVER, THE TIMING OF CONDITIONAL ACTION IS SUCH THAT ALL T-STEPS CONTAINING A CONDITIONAL SKIP OR GATE NS MUST BE STRETCHED. IN ADDITION, THE TEST CONDITION MUST BE STABLE AT THE TIME THE TEST IS MADE. THIS IS ASSURED BY THE RULES BELOW.

THE LOCAL TEST CONDITIONS ARE STABLE AT THE TIME THE OUTPUT OF THE SHIFTER OR ALU IS STABLE FOR GATING INTO LOCAL STORE. THESE TIMES ARE COVERED IN SECTION 5.6.3. CONDITIONAL ACTION BASED ON THE LOCAL TEST CONDITIONS MAY OCCUR IN THE T-STEP IMMEDIATELY FOLLOWING THE ONE IN WHICH THE TRAILING-EDGE FUNCTIONS ARE VALID. ALTERNATIVELY, CONDITIONAL ACTION IS VALID IN A STRETCHED T-STEP WHEN THE GATE FUNCTIONS WOULD HAVE BEEN VALID IN THE SAME UNSTRETCHED T-STEP.

THE GLOBAL TEST CONDITIONS ARE THE SIX BITS OF FIST. THIS REGISTER CAN BE LOADED WITH AN ORDINARY SIX-BIT TRANSFER OPERATION, OR AS FOLLOWS:

IF THE "ALU STATUS ENABLE" BIT IS ON IN THE ACTIVE K-VECTOR, THEN WHEN GATE ALU IS EXECUTED:

- A) THE LOCAL CONDITIONS S, R, AND O ARE LOADED INTO THEIR COUNTERPARTS IN FIST (TRAILING-EDGE).
- B) THE "C" BIT IN FIST IS LOADED AS FOLLOWS (TRAILING-EDGE):
 - 1. IF A NANOPRIMITIVE IS SIMULTANEOUSLY EXECUTED TO LOAD COH, THE VALUE LOADED INTO COH IS ALSO LOADED INTO FIST-C.
 - 2. IF NO SUCH OPERATION IS SIMULTANEOUSLY COMMANDED, FIST-C IS LOADED FROM THE CURRENT VALUE OF COH (THE CURRENT LOCAL CARRY CONDITION).

IF THE "SH STATUS ENABLE" BIT IS ON IN THE ACTIVE K-VECTOR, THE TWO LOCAL CONDITIONS SHB AND SLB ARE LOADED INTO THE TWO CORRESPONDING BITS OF FIST WHEN GATE SH IS EXECUTED (TRAILING-EDGE).

CONDITIONAL ACTION BASED ON THE GLOBAL TEST CONDITIONS IS VALID IN THE T-STEP IMMEDIATELY FOLLOWING A T-STEP IN WHICH THE STATE OF FIST IS CHANGED.

THE SPECIAL CONDITIONS MS BUSY AND MS DATA INVALID ARE SYNCHRONIZED BY THE MACHINE CLOCK. THUS THEY MAY BE TESTED IN ANY STRETCHED T-STEP. SECTION 5.4.3 COVERS MAIN STORE TIMING.

THE SPECIAL CONDITION F NOT ZERO REFLECTS THE STATE OF THE F REGISTER SPECIFIED BY FSEL1 AT THE LEADING EDGE OF THE TESTING T-STEP. THE SPECIAL CONDITION R INDEX NOT ZERO REFLECTS THE RESULT OF THE LAST INDEX ALU OPERATION PRIOR TO

THE LEADING EDGE OF THE TESTING T-STEP. THE SPECIAL CONDITION PROGRAM CHECK REFLECTS THE CURRENT COMBINED STATE OF ALL THE PROGRAM CHECK CONDITIONS IN EXTERNAL STORE REGISTERS 20 AND 21. (SEE SECTION 4.5.2.2).

NOTE: UNCONDITIONAL SKIP, GATE NS ARE SIMPLE TRAILING-EDGE OPERATIONS. THEY DO NOT REQUIRE THAT THE T-STEP BE STRETCHED.

5.8 MISCELLANEOUS FUNCTIONS

5.8.1 GENERATE INTERRUPT

IN ORDER TO FACILITATE USING THE EXTERNAL INTERRUPT STRUCTURE FROM WITHIN A NANOPROGRAM, THE GENERATE INTERRUPT COMMAND IS PROVIDED. THIS COMMAND ALLOWS THE NANOPROGRAMMER TO GENERATE OR CLEAR ANY INTERRUPT INTERNALLY.

THE VALUE OF G(GSPEC) IN T-PERIOD 1 SELECTS THE INTERRUPT LEVEL TO BE LATCHED OR UNLATCHED WHEN THE GENERATE INTERRUPT K-BIT IS SET (SEE SECTION 4.5.2.4).

WHEN THE HIGH-ORDER BIT OF G(GSPEC) IS 1, THE LOW-ORDER FIVE BITS OF G(GSPEC) SELECT THE INTERRUPT LEVEL (2-31) TO BE CLEARED (VALUES 0 AND 1 SPECIFY NO ACTION).

WHEN THE HIGH-ORDER BIT OF G(GSPEC) IS 0, THE LOW-ORDER FIVE BITS OF G(GSPEC) SELECT THE INTERRUPT LEVEL (2-31) ON WHICH TO GENERATE AN INTERRUPT (VALUES 0 AND 1 SPECIFY NO ACTION).

5.8.2 AUXILIARY ACTION

THE AUXILIARY ACTION NANOPRIMITIVE USES THE VALUE FOUND IN F REGISTER FACT TO COMMAND SPECIAL QM-1 CONTROL FUNCTIONS. IN GENERAL, THESE AUXILIARY COMMANDS WILL ENABLE OR DISABLE VARIOUS INTERRUPT AND CONTROL FACILITIES. THE FORMAT OF THE SIX-BIT CONTROL FIELD IS AS FOLLOWS:

```

-----
FACT / BIT 5  BIT 4  BIT 3  BIT 2  BIT 1  BIT 0 /
-----

```

CURRENTLY DEFINED AUXILIARY ACTION COMMANDS ARE:

COMMAND (OCTAL)	FUNCTION
00	NO OPERATION. (ALL UNDEFINED COMMANDS ARE ALSO TREATED AS NO OPERATIONS.)
77	DISABLE INTERRUPT LEVELS 2-31. OVERRIDES THE ALLOW INTERRUPT BITS IN THE ACTIVE, AND ALL SUBSEQUENT, K-VECTORS, UNTIL RESCINDED BY THE ENABLE INTERRUPT COMMAND BELOW.
76	ENABLES INTERRUPTS. RESCINDS THE ACTION OF THE DISABLE INTERRUPT COMMAND ABOVE, AND RESTORES THE CONTROL OF INTERRUPTS TO THE ALLOW INTERRUPT BITS. NOTE: ONLY THOSE INTERRUPT LEVELS MASKED ON" BY THE INTERRUPT ENABLE BITS IN EXTERNAL STORE REGISTERS 18 AND 19 ARE ENABLED.

BOTH THE IO INTERRUPT DISABLE AND ENABLE COMMANDS BECOME EFFECTIVE WITHIN 2 T-PERIODS AFTER EXECUTION OF THE T-STEP CONTAINING THE AUXILIARY ACTION NANOPRIMITIVE. IF THE NANOWORD CONTAINING THIS COMMAND ALLOWS INTERRUPTS, IT MUST NOT EXECUTE A READ NS USING OTHER THAN THE PRIORITY BRANCH ADDRESS DURING THESE 2 T-PERIODS. THE RESULT OF A PREMATURE READ NS IS UNDEFINED.

5.8.3 MISCELLANEOUS CONTROL FUNCTIONS

THE FOLLOWING EXTERNAL CONTROL BUTTONS AND SWITCHES APPLY TO THE QM-1 CPU:

MASTER CLEAR - PUSH BUTTON FOR SETTING THE INITIAL CONDITIONS FOR MACHINE START-UP.

START - PUSH BUTTON FOR INITIATING CPU OPERATION

SINGLE/RUN/DOUBLE - A THREE POSITION SWITCH WHICH SELECTS EITHER A SINGLE T-STEP, DOUBLE T-STEP, OR CONTINUOUS OPERATION MODE WHEN THE START BUTTON IS DEPRESSED.

MICRO-STEP - A TWO POSITION SWITCH THAT CAUSES CONTINUOUS OPERATION TO BE STOPPED FOLLOWING EACH EXECUTION OF THE LOAD R31 COMMAND.

5.9 REFERENCE LAYOUT AND MAPS

5.9.1 CPU REGISTER ASSIGNMENT AND LAYOUT

LOCAL STORE	EXTERNAL STORE	F STORE
R0	E0 PORT 0	F0 FMIX
R1	E1 PORT 1	F1 FMOD
R2	E2 PORT 2	F2 FCIA
R3	E3 PORT 3	F3 FAIL
R4	E4 PORT 4	F4 FCID
R5	E5 PORT 5	F5 FAIR
R6	E6 PORT 6	F6 FCOD
R7	E7 PORT 7	F7 FRAD
R8	E8 INDEX 0	F8 FSID
R9	E9 INDEX 1	F9 FSOD
R10	E10 INDEX 2	F10 FEID
R11	E11 INDEX 3	F11 FEOD
R12	E12 INDEX 4	F12 FEIA
R13	E13 INDEX 5	F13 FEOR
R14	E14 INDEX 6	F14 FACT
R15	E15 INDEX 7	F15 FLIV
R16	E16 INDEX 8, MS BASE ADR.	F16 FMPC
R17	E17 INDEX 9, MS FIELD LENGTH	F17 FIDX
R18	E18 INDEX 10, INTERRUPT ENA.	F18 FIST
R19	E19 INDEX 11, BITS	F19 FIPH
R20	E20 INTERRUPT PEND.	F20 G0
R21	E21 FLAGS	F21 G1
R22	E22 INT ADR (2-4)	F22 G2
R23	E23 (5-7)	F23 G3
R24 MPC REG	E24 (8-10)	F24 G4
R25 MPC REG	E25 (11-13)	F25 G5
R26 MPC REG	E26 (14-16)	F26 G6
R27 MPC REG	E27 (17-19)	F27 G7
R28	E28 (20-22)	F28 G8
R29	E29 (23-25)	F29 G9
R30	E30 (26-28)	F30 G10
MIR R31 "C", "A", "B"	E31 (29-31)	F31 G11

THE INDICATOR LAMP LAYOUT FOR ALL CPU REGISTERS IS SHOW IN FIGURE 5.9.1A.

5.9.2 NANOSTORE MAP

K SEGMENT LAYOUT

BYTE	17	16	15	14	13	12	11	10	09	08	07	06	05	04	03	02	01	00	BITS
00	ILE/BR/AN/AM/DA/H1/AS/SS/												KN				I	17-00	
01	ISM/GI/AL/**/**/H2/								KA		/		KB				I	35-18	
02	I			KSHC			/		KALC			/		KSHA				I	53-36
03	I			KS			/		KX			/		KT				I	71-54

LE	LEGAL MICRO ENTRY	H1	HOLD	AS	ALU STATUS ENABLE
SM	SUPERVISOR	H2	HOLD 2	SS	SH STATUS ENABLE
BR	BRANCH	AM	ALLOW MICRO INTERRUPT	GI	GENERATE INTERRUPT
AL	ALTERNATE	AN	ALLOW NANO INTERRUPT	DA	DIRECT MS ACCESS

I SEGMENT LAYOUT

T1	T2	T3	T4	17	16	15	14	13	12	11	10	09	08	07	06	05	04	03	02	01	00	BITS
04	08	12	16	IMG/MR/GM/ RMI /RN/NN/IM/I0/										AUX0		/00/	FSEL0					I 17-00
05	09	13	17	IXI/RI/LE/GE/**/GU/31/GS/I1/										AUX1		/01/	FSEL1					I 35-18
06	10	14	18	ICARR CTL/TEST SPC/TA/GA/I2/										AUX2		/02/	FSEL2					I 53-36
07	11	15	19	IRC/NC/GC/CS A SEL/LNPC /ST/										AUX3		/AA/IN/	GSPEC					I 71-54

MO - MSGO	RC - READ CS	LE - LOAD ES	CARR CTL - CARRY CTL
MR - MSRS	NC - WRITE CS	GE - GATE ES	TEST SPEC- TEST SPECIFIER
GM - GATE MS	GC - GATE CS	GS - GATE SH	CS A SEL - CS ADDR SELECT
RMI- RMI SELECT	IM - INC MPC	GA - GATE ALU	LNPC - LOAD NPC
RN - READ NS	XI - XIO	31 - LOAD R31	GU - GATE NS NCONDITIONLLY
NN - WRITE NS	RI - RIO	TA - TEST ACTION	I0, I1, I2 - IN0, IN1, IN2
ST - STRETCH	IN - INDEX	AA - AUX ACTION	00, 01, 02 - OUT0, OUT1, OUT2

6. NANOPROGRAMMING LANGUAGE SPECIFICATION

6.1 GENERAL

NANOPROGRAMMING IS THE PROCESS OF DEFINING QM-1 HARDWARE CONTROL SEQUENCES AND IMPLEMENTING THEIR DEFINITIONS BY PROGRAMMING THE CONTENTS OF WORDS IN QM-1 NANOSTORE. A NANOPROGRAM, OR LOGICALLY COMPLETE CONTROL SEQUENCE, CAN BE INVOKED FROM ONE OF THREE SOURCES: MACHINE START, INTERRUPT ENTRY, OR MICROINSTRUCTION ENTRY. SINCE IT IS PLAUSIBLE TO REGARD MICROINSTRUCTION CONTROL AS THE "TYPICAL" MODE OF QM-1 OPERATION, MOST NANOPROGRAMMING CAN BE CONSIDERED TO BE THE PROCESS OF DEFINING MICROINSTRUCTIONS AND IMPLEMENTING THEIR DEFINITIONS BY PROGRAMMING THE APPROPRIATE NANOPRIMITIVES IN SEQUENCES OF NANOSTORE WORDS. SUCH A SEQUENCE OF NANOWORDS IS CALLED THE "NANOPROGRAM" CORRESPONDING TO THE DEFINED "MICROINSTRUCTION".

WHILE NANOPROGRAMMING IS THE MOST ELEMENTARY LEVEL OF PROGRAMMING POSSIBLE IN THE QM-1 AND HAS MANY UNIQUE CHARACTERISTICS DEPENDENT ON QM-1 HARDWARE, IT HAS MUCH IN COMMON WITH ANY TYPE OF PROGRAMMING. IN PARTICULAR, IT SHARES THE NEED FOR A SYMBOLIC LANGUAGE TO RELIEVE THE PROGRAMMER OF HAVING TO REMEMBER THE DETAILS OF ACTUAL BIT LOCATIONS AND ABSOLUTE CODES. THIS SECTION DEFINES A NANOPROGRAMMING LANGUAGE TO MEET THE NEEDS OF THE NANOPROGRAMMER, MUCH AS AN ASSEMBLY LANGUAGE MEETS THE NEEDS OF PROGRAMMERS ON A MORE CONVENTIONAL COMPUTER.

ALTHOUGH SOME IMPLEMENTATION STANDARDS ARE DESCRIBED, THE INTENT OF THIS SECTION IS TO PRESENT A LANGUAGE SPECIFICATION ONLY, GENERALLY AVOIDING ELEMENTS (SUCH AS ASSEMBLER DIRECTIVES) THAT MORE PROPERLY BELONG IN AN ASSEMBLER USER'S MANUAL.

WE WOULD LIKE TO ACKNOWLEDGE THE EFFORTS OF DR. BOB NASH, AT THE DEPARTMENT OF COMPUTER SCIENCE, STATE UNIVERSITY OF NEW YORK AT BUFFALO. THE NANO-ASSEMBLER IS BASED ON HIS DEFINITION OF THE "NANOCODE SYMBOLIC ASSEMBLER", DEVELOPED AT THE UNIVERSITY.

1. NOTES APPEARING WITHIN THIS CHAPTER, ENCLOSED IN BRACKETS, DENOTE TEMPORARY RESTRICTIONS OR FEATURES FOUND ONLY IN THE BASIC NANO-ASSEMBLER (VERSION 1; LEVEL 2). J

6.2 ELEMENTS

6.2.1 SOURCE STATEMENTS

SOURCE STATEMENTS CONSIST OF SINGLE RECORDS THAT ARE EITHER:

- A) COMMENT STATEMENTS
- B) LABEL STATEMENTS
- C) COMMAND STATEMENTS
- D) CONTROL STATEMENTS (BRIEFLY DISCUSSED IN THIS DOCUMENT)

EACH TYPE OF STATEMENT IS COMPLETE ON A SINGLE RECORD. THERE IS NO PROVISION FOR CONTINUATION OF A STATEMENT.

SETS OF COMMAND STATEMENTS ACTUALLY DEFINE NANOWORDS. THESE SETS ARE PRECEDED BY A LABEL STATEMENT TO LOCATE AND NAME THE NANOWORD SO DEFINED.

COMMENT STATEMENTS ARE USED ONLY TO ANNOTATE THE LISTING AND ARE OTHERWISE IGNORED. CONTROL STATEMENTS PROVIDE INFORMATION AND DIRECTION FOR THE PROCESSOR THAT TRANSLATES THE NANOPROGRAMS INTO ABSOLUTE BIT STRINGS IN A FORM SUITABLE FOR LOADING AND EXECUTION ON THE QM-1.

ALL STATEMENTS, EXCLUDING COMMENTS, MAY BE SUBDIVIDED INTO FIELDS. ANY NUMBER OF FIELDS MAY OCCUR ON LABEL, COMMAND, AND CONTROL STATEMENTS, AS WARRANTED BY THEIR IMMEDIATE APPLICATIONS. FIELDS ARE SEPARATED BY FIELD DELIMITERS, AS DESCRIBED BELOW (6.2.4.1).

EACH TYPE OF STATEMENT WILL BE DEFINED IN MORE DETAIL IN SUBSEQUENT SECTIONS.

6.2.2 CHARACTER SET

THE CHARACTER SET AVAILABLE FOR WRITING NANOPROGRAMS IS A SUBSET OF BOTH ASCII AND THE IBM 829 KEYPUNCH. IT INCLUDES THE FOLLOWING CHARACTERS:

- | | |
|--------------|---------------------|
| A, B, ..., Z | UPPER CASE ALPHABET |
| 0, 1, ..., 9 | DECIMAL DIGITS |
| " " | BLANK |
| . | PERIOD |
| , | COMMA |
| : | COLON |
| = | EQUAL |

>	GREATER-THAN-SIGN
+	PLUS
-	MINUS
/	SLASH
*	STAR
"	QUOTATION MARK

THIS CHARACTER SET WILL BE EXPANDED AS THE NEED FOR ADDITIONAL SPECIAL CHARACTERS ARISES.

6.2.3 SYMBOLIC NAMES

SYMBOLIC NAMES ARE STRINGS OF LETTERS, DIGITS, PERIODS, AND SINGLE OCCURRENCES OF THE BLANK CHARACTER PRECEDED BY ANY NON-BLANK CHARACTER. A SYMBOLIC NAME MAY BEGIN WITH A LETTER OR PERIOD. LEADING BLANKS ARE IGNORED. THESE STRINGS MAY BE OF ANY LENGTH, BUT ONLY THE FIRST 10 CHARACTERS ARE USED FOR RECOGNITION OF SYMBOLIC NAMES. IF THE 10-TH CHARACTER OF A NAME IS A BLANK IT IS ALSO IGNORED.

6.2.4 DELIMITERS

6.2.4.1 FIELD DELIMITERS

LABEL, COMMAND, AND CONTROL STATEMENTS MAY, OPTIONALLY, BE DIVIDED INTO FIELDS. TWO FIELD DELIMITERS ARE DEFINED. THE COMMA (,) IS MOST FREQUENTLY USED TO SEPARATE FIELDS, AND IS TREATED AS A FIELD DELIMITER ON ALL STATEMENT TYPES. IT IS IGNORED WITHIN COMMENT FIELDS (6.2.5.2) AND ON COMMENT STATEMENTS. THE BLANK () MAY ALSO BE USED AS A FIELD DELIMITER, BUT ONLY WHEN TWO OR MORE BLANKS IMMEDIATELY FOLLOW A LEGAL SYMBOLIC NAME, AND WHERE THE FOLLOWING FIELD BEGINS WITH A SYMBOLIC NAME. THE REASONS FOR THIS ALTERNATE DELIMITER ARE DISCUSSED BELOW, IN THE SECTION ON PSEUDO COMMANDS (6.3.2.1).

1. VERSION 1, LEVEL 2 RESTRICTION: ANY OCCURENCE OF TWO OR MORE BLANKS FOLLOWING A SYMBOLIC NAME WILL BE TREATED AS A FIELD DELIMITER, REGARDLESS OF THE FIRST COMPONENT OF THE FOLLOWING FIELD.]

6.2.4.2 SYMBOLIC NAME DELIMITERS (OPERATORS)

SYMBOLIC NAMES MAY BE DELIMITED BY EITHER FIELD DELIMITERS (6.2.4.1) OR OPERATORS. THE BASIC ARITHMETIC OPERATORS + (ADDITION), - (SUBTRACTION), * (MULTIPLICATION), / (DIVISION), AND = (ASSIGNMENT) ARE RECOGNIZED ONLY WITHIN COMMAND OR CONTROL FIELDS WHERE ARITHMETIC EXPRESSIONS ARE LEGAL. ONE ADDITIONAL OPERATOR IS DEFINED, AND CONSISTS OF THE CHARACTER PAIR "->" (TRANSMIT). THIS OPERATOR MAY BE USED ONLY WHERE DATA TRANSFER COMMANDS ARE LEGAL. FINALLY, THE QUOTATION MARK (") MAY FOLLOW A SYMBOLIC NAME ACTING AS BOTH A COMMENT FIELD DELIMITER (SEE 6.2.5.2 BELOW) AND A SYMBOLIC NAME DELIMITER.

6.2.5 COMMENTS

6.2.5.1 COMMENT STATEMENTS

ANY STATEMENT THAT HAS * IN COLUMN ONE OF THE STATEMENT WILL BE TREATED AS A COMMENT STATEMENT. IT WILL BE PRINTED ON THE SOURCE LISTING BUT WILL HAVE NO OTHER EFFECT ON THE TRANSLATION.

6.2.5.2 COMMENT FIELDS

COMMENTS MAY BE INCLUDED WITHIN FIELDS ON LABEL, COMMAND, AND CONTROL STATEMENTS BY SIMPLY ENCLOSING THE COMMENT BETWEEN A PAIR OF QUOTATION MARKS ("). COMMENT FIELDS MAY BE PLACED BEFORE OR AFTER SYMBOLIC NAMES. A COMMENT FIELD THAT IS IN EFFECT IS TERMINATED UPON ENCOUNTERING THE END OF A STATEMENT. A NEW COMMENT FIELD MUST BE DECLARED ON THE NEXT STATEMENT, IN ORDER TO CONTINUE THAT COMMENT.

6.2.6 BLANKS

ASIDE FROM THEIR USE IN SYMBOLIC NAMES (6.2.3) AND AS FIELD DELIMITERS (6.2.4.1) STRINGS OF BLANKS ARE IGNORED.

6.3 NANOWORD DEFINITION

6.3.1 LABEL STATEMENT

A LABEL STATEMENT IS DEFINED AS A STATEMENT THAT CONTAINS EITHER A SYMBOLIC OR NULL LABEL DECLARATION.

A SYMBOLIC LABEL IS INDICATED BY A FIELD CONTAINING A SYMBOLIC NAME FOLLOWED BY A COLON.

A NULL LABEL CONSISTS OF A FIELD CONTAINING ONLY A COLON, WITH NO PRECEDING CHARACTERS.

THE COLON TERMINATES THE LABEL FIELD, BUT NOT THE LABEL STATEMENT. ADDITIONAL FIELDS CONTAINING PSEUDO COMMANDS (SEE SECTION 6.3.2.1) MAY BE INCLUDED ON THE STATEMENT.

A SYMBOLIC LABEL IS USED SPECIFICALLY TO PASS THE SYMBOLIC NAME, AND CORRESPONDING NANOSTORE LOCATION, OF A MICRO-INSTRUCTION TO THE MICRO-ASSEMBLER (DESCRIBED IN ANOTHER DOCUMENT); FOR USE AS AN ACTUAL OPERATION CODE.

I IN VERSION 1, LEVEL 2, THIS LABEL CANNOT BE USED FOR REFERENCE BY ANY OTHER NANOWORD; INSTEAD A VARIABLE SYMBOL NAME MAY BE EQUIVALENCED TO THE SAME NANOSTORE LOCATION (SEE EXAMPLE BELOW).]

THE OCCURRENCE OF A LABEL FIELD INDICATES THE END OF THE PRECEDING NANOWORD (IF ANY) AND BEGINS A NEW WORD DEFINITION.

EXAMPLES OF LABEL STATEMENTS:

1. BEGINNING OF AN ADD INSTRUCTION.
ADD: "MICRO INSTRUCTION FORMAT ATTRIBUTES GO HERE"
2. COMPLEX OPERATION CODE NAME AND ATTRIBUTES.
DECODE X.R: MICRO = ABSOLUTE + AB RELATIVE + WORD 2
("MICRO" IS EXPLAINED IN SECTION 6.3.2.1)
3. VARIABLE NAME USED FOR REFERENCE TO THIS NANOWORD; THROUGH KN FIELDS OF OTHER NANOWORDS. THE SPECIAL SYMBOLIC NAME "N." IS USED TO ACCESS THE CURRENT NANOWORD ADDRESS (SEE SECTION 6.3.2.1).
ADD: ADD. OPR = N., MICRO = A B ABSOLUTE.

[VERSION 1, LEVEL 2, MAY USE ONLY THE METHOD IN EXAMPLE 3 TO
REFERENCE OTHER NANOWORDS.]

6.3.2 COMMAND STATEMENTS

COMMAND STATEMENTS SERVE TO ACTUALLY DEFINE THE NANOPRIMITIVES DESIRED IN A NANOWORD, AND TO SHOW IN WHICH T-VECTOR THEY SHOULD APPEAR. EACH STATEMENT CONSISTS OF ONE OR MORE COMMAND FIELDS, WHICH ARE ORDER INDEPENDENT WITHIN EACH T-VECTOR [NOT SO IN VERSION 1, LEVEL 2; SEE SECTION 6.3.2.4].

THERE ARE THREE CLASSES OF COMMANDS: PSEUDO COMMAND OPERATORS, NANOPRIMITIVE COMMANDS, AND 6 BIT DATA TRANSFER COMMANDS. PSEUDO COMMAND OPERATORS ARE USED TO DECLARE ASSEMBLY TIME FUNCTIONS AND ATTRIBUTES, AFFECTING THE NANOWORD CURRENTLY BEING ASSEMBLED. NANOPRIMITIVE COMMANDS EACH DEFINE VALUES FOR UP TO TWO K-VECTOR OR T-VECTOR FIELDS. THESE COMMANDS EXPLICITLY IDENTIFY NANOWORD FIELDS TO BE INCLUDED IN THE GENERATED NANOPROGRAM, BINARY OUTPUT FILE. THE 6 BIT DATA TRANSFER COMMANDS ARE EFFECTIVELY MACROSCOPIC NANOPRIMITIVE COMMANDS, AND DEFINE NANOWORD FIELDS AFFECTING TRANSFERS BETWEEN 6 BIT AUX FIELDS AND F REGISTERS. A SINGLE TRANSFER COMMAND MAY IMPLICITLY DEFINE UP TO 5 NANOWORD FIELDS. ALL THREE COMMAND CLASSES ARE DISCUSSED IN MORE DETAIL IN THE FOLLOWING SECTIONS.

6.3.2.1 PSEUDO COMMAND OPERATORS

PSEUDO COMMAND OPERATORS ARE REQUIRED TO SPECIFY ASSEMBLY TIME ATTRIBUTES OF THE NANOWORD CURRENTLY BEING PROCESSED. THESE ATTRIBUTES INCLUDE EXTERNAL MICRO-INSTRUCTION FORMAT INDICATORS, FOR USE BY THE MICRO-ASSEMBLER, AND SELECTION OF THE APPROPRIATE T-VECTORS TO RECEIVE SPECIFIC NANOPRIMITIVE COMMANDS. THERE ARE THREE TYPES OF PSEUDO COMMANDS SUPPORTED IN THE VERSION 1 NANO-ASSEMBLER: NANOSTORE LOCATION COUNTER, MICRO-INSTRUCTION ATTRIBUTES, AND POSITION DECLARATIONS.

ALL PSEUDO COMMAND OPERATORS MAY BE DELIMITED BY A COMMA, OR AT LEAST TWO BLANK CHARACTERS (AS DESCRIBED IN SECTION 6.2.4.1 ABOVE). USE OF BLANKS AS THE DELIMITER ALLOWS PSEUDO COMMANDS TO TAKE ON THE APPEARANCE OF OPERATORS, AS FOUND IN MORE CONVENTIONAL ASSEMBLERS, WHERE BLANKS DENOTE SEPARATION BETWEEN OPERATOR AND PARAMETER FIELDS.

N A N O S T O R E L O C A T I O N C O U N T E R

THE NANOSTORE LOCATION COUNTER MAY BE ACCESSED, OR MODIFIED, USING ARITHMETIC EXPRESSIONS. THIS CAPABILITY PERMITS THE ASSIGNMENT OF THE CURRENT NANO-LOCATION-COUNTER TO ANOTHER SYMBOLIC NAME. THIS SYMBOL MAY THEN BE REFERENCED THROUGH KN FIELDS OF OTHER NANOWORDS, PERMITTING SYMBOLIC NANO-BRANCH DECLARATIONS. CONTINUATION NANOWORDS, IN THE LARGER NANOPROGRAMS, NEED NOT HAVE SYMBOLIC LABELS (ARE NOT KNOWN TO THE MICRO-ASSEMBLER) AND MUST THEREFORE BE REFERENCED THROUGH THIS ALTERNATE MEANS.

: SEARCH. 2 = N. "LABEL STATEMENT WITH NULL LABEL FIELD"

IT IS LEGAL TO MODIFY THE NANO-LOCATION-COUNTER, AS A METHOD OF ALTERING THE NORMALLY SEQUENTIAL ORDER OF CODE GENERATION. CARE MUST BE TAKEN WHEN USING THIS METHOD SINCE THE VALUE OF N. MUST BE SET TO ONE LESS THAN THAT OF THE NEXT DESIRED NANOWORD LOCATION.

"MAKE THE ADD INSTRUCTION MICRO-OPERATION CODE 40" N. = 40-1
ADD: ADD LOC = N. "ADD LOC = LOCATION OF ADD"

I THE BASIC VERSION 1 ASSEMBLERS HAVE CERTAIN PREDEFINED VARIABLE SYMBOL NAMES, TO SIMPLIFY DECLARATION OF FUNCTIONAL UNIT ACTIONS (SUCH AS ADD, SUB, OVERFLOW, RESULT, ETC.). THE TEMPORARY DIFFERENTIATION BETWEEN SYMBOLIC LABELS AND SYMBOLIC NAMES PERMITS USE OF ALTERNATE NANOWORD REFERENCE NAMES, WHERE THE PREDEFINED SYMBOLIC NAME VALUE MUST BE RETAINED.)

M I C R O - I N S T R U C T I O N A T T R I B U T E D E C L A R A T I O N S

ONLY A LIMITED SET OF MICRO-INSTRUCTION FORMATS ARE PREDEFINED WITHIN THE VERSION 1 MICRO-ASSEMBLER. SELECTION OF THE FORMAT TO BE USED BY THAT ASSEMBLER, WHEN ENCOUNTERING SPECIFIC MICRO-INSTRUCTION OPERATION CODES, IS PROVIDED BY THE "MICRO" PSEUDO COMMAND AT NANO-ASSEMBLY TIME. SELECTIONS MAY, OPTIONALLY, BE MADE DURING MICRO-ASSEMBLIES.

VERSION 1 FORMAT SELECTION IS PROVIDED BY A FORMAT INDEX NUMBER, PASSED BETWEEN THE TWO ASSEMBLERS. THE ACTUAL FORMAT DECLARATION SHOULD BE MADE ON THE LABEL STATEMENT, FOLLOWING THE SYMBOLIC LABEL FIELD. THE "MICRO" PSEUDO COMMAND APPEARS AS FOLLOWS:

MICRO = <EXPRESSION>

WHERE <EXPRESSION> MAY BE ANY LEGAL ARITHMETIC EXPRESSION (AS DISCUSSED IN SECTION 6.3.2.2). THE LIST BELOW DESCRIBES SOME OF THE FORMATS SUPPORTED BY THE VERSION 1 MICRO-ASSEMBLER. A ZERO VALUED MICRO DECLARATION (OR NO DECLARATION) CAUSES THE MICRO-ASSEMBLER TO SELECT A DEFAULT FORMAT.

MICRO VALUE --(OCTAL)--	FORMAT ATTRIBUTES -----
200	"OP M,N" 18 BIT, WITH ABSOLUTE (*) INSTRUCTION PARAMETERS. M IS THE 5 BIT A FIELD, N IS THE 6 BIT B FIELD.
201	"OP M,N" 18 BIT, WITH PARAMETER M AS A 5 BIT ABSOLUTE A FIELD. N IS THE 6 BIT B FIELD, WITH A MICRO-LOCATION-COUNTER RELATIVE VALUE.
207	"OP MN" 18 BIT, WITH PARAMETER MN REPRESENTING AN 11 BIT, MICRO-LOCATION-COUNTER RELATIVE VALUE.
300	"OP M,V,N" 36 BIT, WITH M REPRESENTING THE 5 BIT A FIELD (ABSOLUTE). V REPRESENTS THE 18 BIT SIGNED VALUE OF THE SECOND WORD OF THE INSTRUCTION. N IS THE 6 BIT B FIELD (ABSOLUTE).
303	"OP MN,V" 36 BIT, WITH MN REPRESENTING AN 11 BIT ABSOLUTE VALUE (AB FIELD). V REPRESENTS THE 18 BIT SIGNED VALUE OF THE SECOND WORD.

- * IN THE ABOVE TABLE; LOCATION-COUNTER RELATIVE VALUES ARE ALL SIGNED, TWO'S COMPLEMENT, WITH LEFTMOST BIT INDICATING THE SIGN. THE TERM "ABSOLUTE" REFERS TO NON-RELOCATABLE ADDRESS EXPRESSIONS, AS RELOCATABLE EXPRESSIONS ARE ILLEGAL WITHIN THOSE PARAMETERS.

POSITION DECLARATION COMMANDS

THE POSITION DECLARATION COMMANDS DETERMINE WHETHER THE COMMANDS THAT FOLLOW SPECIFY K-VECTOR OR T-VECTOR FIELDS, AND IN THE CASE OF T-VECTOR SPECIFICATIONS ALSO SELECT THE T-VECTOR POSITION. A POSITION DECLARATION MUST APPEAR PRIOR TO ANY T-VECTOR COMMANDS. K-VECTOR COMMANDS MAY APPEAR ANYWHERE FOLLOWING THE LABEL STATEMENT, ALTHOUGH IT IS RECOMMENDED THAT THEY BE PLACED FOLLOWING THE K-VECTOR POSITION DECLARATION, TO AVOID POSSIBLE PROGRAMMING ERRORS. EACH POSITION DECLARATION CONSISTS OF A SINGLE SYMBOLIC NAME, WHICH MAY BE TERMINATED BY EITHER A COMMA OR MULTIPLE BLANK DELIMITER (SEE SECTION 6.2.4.1, FIELD DELIMITERS). THE ALLOWED FORMS OF THESE DECLARATIONS ARE:

... ONLY K COMMANDS ALLOWED IN STATEMENT.
X... ALL T COMMANDS APPLY TO T-VECTOR 1 OF THE CURRENT NANOWORD.
...X... ALL T COMMANDS APPLY TO T-VECTOR 2 OF THE CURRENT NANOWORD.
...X... ALL T COMMANDS APPLY TO T-VECTOR 3 OF THE CURRENT NANOWORD.
...X... ALL T COMMANDS APPLY TO T-VECTOR 4 OF THE CURRENT NANOWORD.

ANY STATEMENTS NOT CONTAINING A POSITION DECLARATION WILL TAKE ON THE SAME POSITION ATTRIBUTES AS THE MOST RECENTLY ENCOUNTERED POSITION DECLARATION. A LABEL STATEMENT POSITIONS THE NEW NANOWORD IN ITS K-VECTOR.

POSITION DECLARATIONS MAY ALSO SPECIFY THAT THE T-VECTOR BE STRETCHED. IN THE LIST ABOVE, THE APPEARANCE OF AN X INDICATES THAT THE STRETCH NANOPRIMITIVE IS NOT SPECIFIED. IF THE "X" IS REPLACED WITH AN "S", THE STRETCH NANOPRIMITIVE IS SPECIFIED TO BE ACTIVE IN THE DECLARED T-VECTOR.

FOR THE PROTOTYPE SYSTEM ONLY, THE LETTER "X" MAY BE REPLACED WITH THE LETTER "P" SPECIFYING AN AUTOMATIC HARDWARE STRETCH IS BEING ACTIVATED BY ANOTHER NANOPRIMITIVE FUNCTION, AND THEREFORE AN EXPLICIT STRETCH IS REQUIRED ONLY ON PRODUCTION MACHINES.

...S... ALL T COMMANDS APPLY TO T-VECTOR 2, AND THE STRETCH NANOPRIMITIVE IS SELECTED.
...P... ALL T COMMANDS APPLY TO T-VECTOR 4, THE STRETCH NANOPRIMITIVE WILL NOT BE SELECTED ON THE PROTOTYPE QM-1.

6.3.2.2 NANOPRIMITIVE COMMANDS

THE NANO-ASSEMBLER PROVIDES COMMANDS FOR THE SPECIFICATION OF ALL DEFINED NANOPRIMITIVE FIELDS IN QM-1. THESE ARE DIVIDED INTO TWO CLASSES, K COMMANDS AND T COMMANDS, TO CORRESPOND TO THE K-VECTOR AND T-VECTOR PORTIONS OF A NANOWORD. IN GENERAL, EACH NANOPRIMITIVE COMMAND SPECIFIES THE VALUE OF ONE, OR MORE, FIELDS IN A K OR T-VECTOR. K-VECTOR FIELDS AND T-VECTOR FIELDS ARE SUMMARIZED IN SECTIONS 5.3.6 AND 5.3.7, RESPECTIVELY, OF THIS MANUAL.

IN THE VERSION 1 NANO-ASSEMBLER, A NANOPRIMITIVE COMMAND MAY APPEAR IN ONE OF FOUR POSSIBLE FORMATS. THESE ARE:

```
NANOPRIMITIVE NAME  
NANOPRIMITIVE NAME=<PRIMARY FIELD EXPRESSION>  
NANOPRIMITIVE NAME<<SECONDARY FIELD EXPRESSION>>  
NANOPRIMITIVE NAME<<SECONDARY EXPRESSION>>=<PRIMARY EXPRESSION>
```

EXPRESSIONS ARE SUPPORTED IN THE VERSION 1 ASSEMBLER BY A SIMPLE LEFT TO RIGHT, ARITHMETIC EVALUATION. THERE IS NO DEFINED OPERATOR PRECEDENCE. CARE MUST BE TAKEN IN THE CODING OF EXPRESSIONS, SINCE THE VERSION 2 ASSEMBLERS WILL INTRODUCE BASIC MULTIPLICATION AND DIVISION PRECEDENCE. ALL RESULTS ARE IN TWO'S COMPLEMENT, SIGNED FORMAT. VERSION 1 DOES NOT CHECK FOR FIELD VALUE OVERFLOW.

IN MANY CASES NANOPRIMITIVE FUNCTIONS REQUIRE AT LEAST ONE ASSOCIATED NANOPRIMITIVE SPECIFICATION, IN ORDER TO COMPLETELY DESCRIBE THE ACTION TO BE PERFORMED. FOR THIS REASON PROVISION IS MADE FOR THE PRIMARY NANOPRIMITIVE COMMAND NAME TO SPECIFY BOTH ITS OWN FIELD VALUE AND, WHERE NECESSARY, A SECONDARY FIELD VALUE. AN EXAMPLE OF THIS IS:

```
READ CS ( CS ADDRESS SELECT )
```

WHERE "READ CS" SETS THE VALUE 1 INTO THE READ CS FIELD, AND THE VALUE OF THE EXPRESSION IN THE SECONDARY FIELD (PARENTHEZIZED) IS PLACED INTO THE CORRESPONDING CS ADDRESS SELECTION FIELD.

THE VERSION 1 NANO-ASSEMBLER MAINTAINS A SET OF PREDEFINED VARIABLE SYMBOL NAMES, AND VALUES, WITHIN ITS SYMBOL TABLE (FURTHER DISCUSSED IN SECTION 6.3.2.3). THESE PREDEFINED SYMBOLS ARE PROVIDED FOR USE WITHIN NANOPRIMITIVE EXPRESSIONS, AND PROVIDE THE CORRESPONDING VALUES FOR MOST FREQUENTLY USED SYMBOLIC NAMES (SUCH AS F000, G2, B, MPC, ADD, CARRY, ETC.). THE FOLLOWING TABLES DESCRIBE THE SUPPORTED NANOPRIMITIVE COMMANDS, IDENTIFY THEIR SECONDARY FIELDS, AND LIST THOSE PREDEFINED SYMBOLIC NAMES AND VALUES SPECIFICALLY PROVIDED FOR REFERENCE USE IN THOSE FIELDS.

THE K-VECTOR NANOPRIMITIVE COMMANDS MAY APPEAR FOLLOWING ANY K-VECTOR, OR T-VECTOR, POSITION DECLARATION.

COMMAND NAME

FUNCTIONS

=====

=====

<OPERATING STATE CONTROL FIELDS>

LEGAL MICRO OP	- SETS THE LEGAL MICRO OP ENTRY BIT.
ALLOW NANO INTERRUPT	- SETS THE ALLOW NANO INTERRUPT BIT.
ALLOW MICRO INTERRUPT	- SETS THE ALLOW MICRO INTERRUPT BIT.
ALLOW INTS	- SETS BOTH ALLOW INTERRUPT BITS.
DIRECT MS ACCESS	- SETS THE DIRECT MAIN STORE ACCESS BIT.
HOLD	- SETS THE HOLD BIT (KALC, KSHC, KSHA, KS).
HOLD 2	- SETS THE HOLD 2 BIT (KA, KB).
SH STATUS ENABLE	- SETS THE SHIFTER STATUS ENABLE BIT.
ALU STATUS ENABLE	- SETS THE ALU STATUS ENABLE BIT.
SUPERVISOR	- SETS THE SUPERVISORY INSTRUCTION BIT.
GENERATE INTERRUPT	- SETS THE GENERATE / CLEAR INTERRUPT ACTIVATION BIT.

<NANO-BRANCH CONTROL FIELDS>

KN=<EXPRESSION>	- SETS THE KN FIELD TO THE VALUE OF <EXPRESSION>.
BRANCH<EXPRESSION>>	- SETS THE NANOBANCH BIT. ALSO ALLOWS THE OPTIONAL SECONDARY FIELD SPECIFICATION OF THE KN FIELD TO THE VALUE OF <EXPRESSION>.
ALT BRANCH<EXPRESSION>>	- SETS THE ALTERNATE BRANCH CONDITION BIT. ALSO ALLOWS KN SPECIFICATION AS IN BRANCH ABOVE.
PREP BRANCH	- SETS BOTH THE NANOBANCH BIT (BRANCH) AND THE ALTERNATE BRANCH CONDITION BIT (ALT BRANCH).

<6 BIT DATA AND FUNCTION CONTROL FIELDS>

THE FOLLOWING 8 FIELDS MAY EACH REFER TO ANY OF THE PREDEFINED SYMBOLS WHICH ARE SPECIFICALLY ORIENTED TOWARD ONLY ONE OR TWO OF THOSE FIELDS. THIS PERMITS THE PLACEMENT OF VALUES, DESTINED FOR A CONTROL FIELD, INTO TEMPORARY HOLDING FIELDS FOR DYNAMIC TRANSFER DURING PROGRAM EXECUTION.

KA=<EXPRESSION>	- SETS KA TO THE VALUE OF <EXPRESSION>.
KB=<EXPRESSION>	- SETS KB TO THE VALUE OF <EXPRESSION>.
KALC=<EXPRESSION>	- SETS KALC TO THE VALUE OF <EXPRESSION>.
ALU CONTROL REFERENCE SYMBOLS PROVIDED ARE: ADD (11), SUB (6), DBL (14), INCR LEFT (17), DECR LEFT (8), PASS LEFT (37),	

PASS RIGHT (32), DECIMAL (40), AND (36), OR (33), XOR (31),
 NOT LEFT (20), NOT RIGHT (25), ZERO (34), ONES (23).
 KSHC=<EXPRESSION> - SETS KSHC TO THE VALUE OF <EXPRESSION>.
 SHIFTER CONTROL REFERENCE SYMBOLS PROVIDED ARE: LEFT (0), RIGHT (1),
 SINGLE (0), DOUBLE (2), CIRCULAR (0), LOGICAL (4), ARITHMETIC (10),
 RIGHT CTL (20), LEFT CTL (40).
 KS=<EXPRESSION> - SETS KS TO THE VALUE OF <EXPRESSION>.
 REFERENCES SAME AS KT BELOW.
 KT=<EXPRESSION> - SETS KT TO THE VALUE OF <EXPRESSION>.
 ARITHMETIC / SHIFT CONDITION (FIST) TEST MASK REFERENCE SYMBOLS
 PROVIDED ARE: SLB (1), OVERFLOW (2), RESULT (4), SIGN (10),
 CARRY (20), SHB (40).
 KX=<EXPRESSION> - SETS KX TO THE VALUE OF <EXPRESSION>.
 MACHINE STATE TEST CONDITION REFERENCE SYMBOLS ARE: F ZERO (1),
 MS DATA (2), MS BUSY (4), PROGRAM CHECK (10), INDEX ZERO (20).

THE T-VECTOR NANOPRIMITIVE COMMANDS MAY APPEAR ONLY AFTER T-VECTOR POSITION
 DECLARATIONS.

COMMAND NAME
 =====

FUNCTIONS
 =====

(NANOSTORE CONTROL COMMANDS)

READ NS - SETS READ NS BIT.
 WRITE NS - SETS WRITE NS BIT.

THE FOLLOWING 3 COMMANDS MAY ALL REFERENCE THE SAME TEST CONDITIONS.

TEST=<EXPRESSION> - SETS THE TEST SPECIFIER FIELD TO THE VALUE OF
 <EXPRESSION>.
 TEST CONDITION REFERENCE SYMBOLS ARE: S (2), NOT S (3), T (4),
 NOT T (5), X (6), NOT X (7).
 SKIP(<EXPRESSION>) - SETS THE GATE NS / SKIP "ACTION" BIT TO ZERO.
 IF THE SECONDARY PARAMETER IS PROVIDED, SETS
 THE TEST SPECIFIER FIELD TO THE VALUE OF
 <EXPRESSION>. IF NO PARAMETER IS SPECIFIED
 (IE. SKIP,) A DEFAULT VALUE OF 1 IS USED
 (UNCONDITIONAL SKIP).

- GATE NS(<EXPRESSION>) - SETS THE GATE NS / SKIP "ACTION" BIT TO 1. THE SECONDARY PARAMETER IS PROCESSED AS FOR THE SKIP NANOPRIMITIVE, ABOVE.
- GATE NS UNCONDITIONAL - SETS THE GATE NS UNCONDITIONAL BIT (OVERRIDES THE CONDITIONAL GATE NS FUNCTION).
- STRETCH - SETS STRETCH BIT OF THE T-VECTOR. MAY BE USED IN PLACE OF (OR ALONG WITH) THE S TYPE POSITION DECLARATION.
- LOAD NPC(<EXPRESSION>) - SETS THE NPC (NANO PROGRAM COUNTER) CONTROL FIELD TO THE VALUE OF <EXPRESSION>. IF THE SECONDARY FIELD IS NOT SPECIFIED THE VALUE 1 (CS) IS USED.

NPC CONTROL REFERENCE SYMBOLS ARE: CS (1), KN (2), SEQ (3).

(CONTROL STORE ACCESS COMMANDS)

THE FOLLOWING 3 COMMANDS MAY ALL REFERENCE THE SAME CS ADDRESSING NAMES.

- READ CS(<EXPRESSION>) - SETS THE READ CS BIT. THE SECONDARY FIELD SPECIFIES THE CS ADDRESS SELECTION CODE. IF NOT SPECIFIED THE VALUE 0 IS USED (CIA).
- CS ADDRESS SELECTION REFERENCE SYMBOLS ARE: CIA (0), COD (1), MPC (2), SWITCHES (7). SELECTION OF MPC RELATIVE ADDRESSES IS ACCOMPLISHED THROUGH ARITHMETIC EXPRESSION BY: MPC+1 (3), MPC+2 (4), MPC+3 "B" (5), AND MPC+4 "AB" (6).
- WRITE CS(<EXPRESSION>) - SETS THE WRITE CS BIT. THE SECONDARY FIELD IS THE SAME AS IN READ CS, ABOVE.
- CS ADDRESS=<EXPRESSION> - SETS THE CS ADDRESS SELECTION FIELD TO THE VALUE OF <EXPRESSION>. SEE READ CS, ABOVE.
- GATE CS - SETS THE GATE CS BIT.
- INC MPC(<EXPRESSION>) - SETS THE INC MPC BIT. SETS THE G SPEC FIELD TO THE VALUE OF <EXPRESSION>. IF THE SECONDARY FIELD IS NOT SPECIFIED A DEFAULT VALUE 14 (MPC PLUS 1) IS USED.

THERE ARE NO PREDEFINED INC MPC REFERENCE NAMES AT THIS TIME. THE FOLLOWING 4 COMMANDS SET BOTH THE INC MPC BIT AND THE PROPER VALUE INTO THE G SPEC FIELD FOR THE DESIRED FUNCTION.

- MPC PLUS 1 - SETS INC MPC AND G SPEC = 14.
- MPC PLUS 2 - SETS INC MPC AND G SPEC = 15.
- MPC PLUS B - SETS INC MPC AND G SPEC = 16.
- MPC PLUS AB - SETS INC MPC AND G SPEC = 17.

LOAD R31 - SETS THE LOAD R31 BIT.

(MAIN STORE ACCESS COMMANDS)

MSG0 - SETS THE MSG0 BIT.
 FETCH MS - SETS THE MSG0 BIT.
 MSRS - SETS THE MSRS BIT.
 WRITE MS - SETS THE MSRS BIT.
 READ MS - SETS BOTH MSG0 AND MSRS BITS.
 RMI=<EXPRESSION> - SETS THE RMI FIELD TO THE VALUE OF
 <EXPRESSION>.
 GATE MS<<EXPRESSION>> - SETS THE GATE MS BIT. IF THE SECONDARY
 FIELD IS SPECIFIED, SETS THE RMI FIELD TO
 THE VALUE OF <EXPRESSION>.

(18 BIT DATA CONTROL COMMANDS)

LOAD ES - SETS THE LOAD ES BIT.
 GATE ES - SETS THE GATE ES BIT.
 GATE SH - SETS THE GATE SH BIT.
 GATE ALU - SETS THE GATE ALU BIT.
 INDEX<<EXPRESSION>> - SETS THE INDEX ALU GATE BIT. IF THE SECONDARY
 FIELD IS SPECIFIED, THE VALUE OF <EXPRESSION>
 IS PLACED INTO THE FSEL2 FIELD (FOR INDEX ALU
 FUNCTION SELECTION).

INDEX ALU FUNCTION REFERENCE SYMBOLS ARE: SUB (6), ADD (11).

INDEX REG<<EXPRESSION 2>>=<EXPRESSION 1>
 - SETS THE VALUE OF THE PRIMARY FIELD
 <EXPRESSION 1> INTO THE AUX2 FIELD (LOCAL STORE
 REGISTER SELECTION) AND THE VALUE OF THE
 SECONDARY FIELD <EXPRESSION 2> INTO THE AUX3
 (INDEX REGISTER SELECTION). EITHER ONE OR
 BOTH PARAMETERS MAY BE SPECIFIED.
 CARRY CTL=<EXPRESSION> - SETS THE CARRY CONTROL FIELD TO THE VALUE OF
 <EXPRESSION>. SPECIFICATION OF THE DESIRED
 CARRY CONTROL FUNCTIONS MAY BE ACCOMPLISHED
 WITH ANY OF THE 7 FOLLOWING COMMANDS.
 CLEAR CIH - SETS THE CARRY CTL FIELD TO THE VALUE 1.
 SET CIH - " " " " " " " " 2.
 ALU TO BOTH CIH AND COH - " " " " " " " " 3.
 ALU TO COH - " " " " " " " " 4.

SET COH	-	"	"	"	"	"	"	"	5.
CLEAR COH	-	"	"	"	"	"	"	"	6.
SH TO COH	-	"	"	"	"	"	"	"	7.

<6 BIT DATA CONTROL COMMANDS>

IN0 - SETS THE IN0 BIT (AUX TO F REGISTER).
 OUT0 - SETS THE OUT0 BIT (F REGISTER TO AUX).
 AUX0<<EXPRESSION 2>>=<EXPRESSION 1>
 - SETS THE AUX0 FIELD TO THE VALUE OF
 <EXPRESSION 1>. IF THE SECONDARY FIELD IS
 SPECIFIED THE IN0 FIELD IS SET TO THE VALUE
 OF <EXPRESSION 2>. THE IN0 FIELD MAY CONTAIN
 ONLY A 0 OR 1.
 FSEL0<<EXPRESSION 2>>=<EXPRESSION 1>
 - SETS THE FSEL0 FIELD TO THE VALUE OF
 <EXPRESSION 1>. IF THE SECONDARY FIELD IS
 SPECIFIED THE OUT0 FIELD IS SET TO THE VALUE
 OF <EXPRESSION 2>. THE OUT0 FIELD MAY CONTAIN
 ONLY A 0 OR 1.

IN1 - SETS THE IN1 BIT.
 OUT1 - SETS THE OUT1 BIT.
 AUX1<<EXPRESSION 2>>=<EXPRESSION 1>
 - SETS THE AUX1 AND IN1 FIELDS AS IN AUX0 ABOVE.
 FSEL1<<EXPRESSION 2>>=<EXPRESSION 1>
 - SETS THE FSEL1 AND OUT1 FIELDS AS IN FSEL0
 ABOVE.

F<<EXPRESSION>>
 - SETS THE FSEL1 FIELD TO THE VALUE OF
 <EXPRESSION>. THIS COMMAND FORMAT IS PROVIDED
 FOR SPECIFICATION OF THE F REGISTER TO BE
 TESTED UNDER THE CONDITIONAL GATE NS / SKIP ON
 "F ZERO" (KX TEST CONDITION).

IN2 - SETS THE IN2 BIT.
 OUT2 - SETS THE OUT2 BIT.
 AUX2<<EXPRESSION 2>>=<EXPRESSION 1>
 - SETS THE AUX2 AND IN2 FIELDS AS IN AUX0 ABOVE.
 FSEL2<<EXPRESSION 2>>=<EXPRESSION 1>
 - SETS THE FSEL2 AND OUT2 FIELDS AS IN FSEL0
 ABOVE.

AUX3=<EXPRESSION>
 - SETS THE AUX3 FIELD TO THE VALUE OF

<EXPRESSION>.

SEVERAL OF THE ABOVE 6 BIT DATA CONTROL FIELDS ARE USED IN SUPPORT OF OTHER FUNCTIONS, SUCH AS INDEX ALU, AND MAY ALSO BE SET AS PRIMARY OR SECONDARY FIELDS OF OTHER NANOPRIMITIVE COMMANDS.

A SET OF PREDEFINED REFERENCE SYMBOLS IS PROVIDED FOR USE BY ALL F SELECT FIELDS. THESE SYMBOLS ARE:

FMIX (0), FMOD (1), FCIA (2), FAIL (3), FCID (4), FAIR (5),
FOOD (6), FROD (7), FSID (10), FSOD (11), FEID (12), FEOD (13),
FEIA (14), FEOR (15), FINV (16), FACT (16), FLIV (17), FMPC (20),
FIDX (21), FIST (22), FIPH (23), G0 (24), G1 (25), G2 (26),
G3 (27), G4 (30), G5 (31), G6 (32), G7 (33), G8 (34),
G9 (35), G10 (36), G11 (37).

A SET OF PREDEFINED REFERENCE SYMBOLS IS PROVIDED FOR SPECIFIC SOURCE AND DESTINATION AUX-SELECT FIELDS. THESE SYMBOLS ARE:

A (0), B (1), C(2) FOR AUX1 SOURCE AND ALL DESTINATIONS, KX (2) FOR
AUX2 SOURCE, KA (3), KB (4) FOR AUX0 AND AUX2 SOURCE AND ALL DEST-
INATIONS, KT (4) FOR AUX1 SOURCE, GSPEC (5) FOR ALL SOURCE,
ALUF (6) FOR AUX0 SOURCE, IO ID (7) FOR AUX0 SOURCE, INCF (6) FOR
AUX1 AND AUX2 SOURCE, DECF (7) FOR AUX1 AND AUX2 SOURCE, KSHC (5)
FOR AUX0 AND AUX3 DESTINATION, KALC (6) FOR AUX0 AND AUX1 DESTIN-
ATION, KSHA (7) FOR AUX1 AND AUX3 DESTINATION, KS. (7) FOR AUX0
DESTINATION, KX. (5) FOR AUX1 DESTINATION, KT. (6) FOR AUX3 DEST-
INATION.

THESE EXPLICIT 6 BIT DATA TRANSFER CONTROLS MAY BE USED AS FOLLOWS:

1. TRANSFER B FIELD (OF R31) TO FSID.
X... AUX0(1)=B, FSEL0=FSID,
2. TRANSFER KX TO KALC (ALU CONTROL) THROUGH FIPH (PHANTOM F).
.S... AUX2(1)=KX, FSEL2(1)=FIPH, AUX3=KALC,

MORE CONCISE COMMAND DECLARATIONS MAY BE OBTAINED USING THE IMPLICIT 6 BIT DATA TRANSFER COMMANDS, AS DESCRIBED IN SECTION 6.3.2.4.

(GENERAL CONTROL AND INPUT / OUTPUT COMMANDS)

- G(<EXPRESSION>) - SETS THE G SPEC FIELD TO THE VALUE OF
<EXPRESSION>. THIS FORMAT IS PROVIDED FOR
SUPPORT OF 6 BIT DATA TRANSFER COMMANDS
(SEE SECTION 6.3.2.4).
- SPECIAL G SPEC SELECTION REFERENCE SYMBOLS ARE: G IO ID (14),
G B (15), G KS (16), G KT (17).
- G SPEC=<EXPRESSION> - SAME AS G(<EXPRESSION>) ABOVE.
- XIO(<EXPRESSION>) - SETS THE XIO BIT. IF THE SECONDARY FIELD IS
SPECIFIED, SETS THE VALUE OF <EXPRESSION> IN
G SPEC.
- RIO(<EXPRESSION>) - SETS THE RIO BIT. IF THE SECONDARY FIELD IS
SPECIFIED, SETS THE VALUE OF <EXPRESSION> IN
G SPEC.
- AUX ACTION - SETS THE AUX ACTION BIT. (ACTIVATES EXTERNAL
COMMANDS VIA USE OF F REGISTER "FACT").

6.3.2.3 VARIABLE SYMBOLS AND CONSTANTS

A VARIABLE SYMBOL TABLE IS PROVIDED, IN THE NANO-ASSEMBLER, TO PERMIT USER DEFINITION OF THEIR OWN SYMBOLIC REPRESENTATIONS FOR MOST REFERENCE SYMBOLS (SEE SECTION 6.3.2.2 FOR LISTS OF REFERENCE SYMBOLS). IN ADDITION MANY PARAMETERS MAY BE SPECIFIED USING USER SELECTED VARIABLE NAMES IN PLACE OF CONSTANTS, PERMITTING EASY MODIFICATION OF SOURCE PROGRAMS. A USER VARIABLE MAY BE REDEFINED AT ANY POINT IN THE ASSEMBLY. SPECIFICATION IS ACCOMPLISHED THROUGH ARITHMETIC ASSIGNMENT STATEMENTS.

THE VARIABLE SYMBOL TO BE DEFINED (OR REDEFINED) APPEARS AS A SYMBOLIC NAME, TO THE LEFT OF AN EQUAL (=) SIGN. THIS NAME MUST NOT BE THE SAME AS ANY CURRENTLY DEFINED NANOPRIMITIVE COMMAND NAME. THE VALUE TO BE PLACED INTO THE VARIABLE IS COMPUTED FROM A SIMPLE ARITHMETIC EXPRESSION, TO THE RIGHT OF THE EQUAL SIGN. THE VERSION 1 NANO-ASSEMBLER SUPPORTS FOUR ARITHMETIC OPERATORS. ALL MULTIPLICATION (*) AND DIVISION (/) OPERATIONS ARE PERFORMED BEFORE ANY ADDITIONS (+) OR SUBTRACTIONS (-). [VERSION 1, LEVEL 2, RECOGNIZES NO OPERATOR PRECEDENCE.] EXPRESSIONS MAY CONSIST OF ANY OTHER VARIABLE NAMES AND CONSTANTS.

CONSTANT VALUES MAY BE SPECIFIED IN OCTAL OR DECIMAL NOTATION. A DECIMAL CONSTANT CONSISTS OF A STRING OF DIGITS FOLLOWED BY A DECIMAL POINT (.). AN OCTAL CONSTANT CONSISTS OF A STRING OF OCTAL DIGITS (0 THROUGH 7) FOLLOWED BY AN OPERATOR OR DELIMITER. TRAILING BLANKS ARE IGNORED.

DECIMAL CONSTANTS: 1., 20., -30790., 0.

OCTAL CONSTANTS: 1, 20, 37707, -77077, 0

ALL NEGATIVE VALUES ARE REPRESENTED IN TWO'S COMPLEMENT NOTATION. CARE MUST BE TAKEN IN SETTING UP MASK VALUES, SINCE -0770 IS NOT THE TRUE COMPLEMENT OF THE VALUE 7007 (-0770 APPEARS AS 777010 IN 10 BIT SIGNED NOTATION). THE VERSION 1 ASSEMBLER SUPPORTS NUMERIC VALUES BETWEEN +32,767 AND -32,768.

VARIABLE ASSIGNMENT EXPRESSIONS MAY BE PLACED WITHIN THEIR OWN FIELDS ON ANY TYPE OF STATEMENT, EXCLUDING COMMENT STATEMENTS, OR ON THEIR OWN STATEMENTS. THE FOLLOWING ARE SOME EXAMPLES OF VARIABLE SYMBOL USE.

1. VARIABLE NAME FOR USE IN CS ADDRESS SELECTION:

S... B FIELD = 3, READ CS (MPC + B FIELD), GATE CS

2. SETTING UP ALTERNATE NAMES FOR SPECIAL PURPOSE F OR G REGISTER REFERENCE:

F.ZERO = 00 "SOURCE OF CONSTANT 6 BIT ZERO AS AN F REGISTER SELECT"

G.ZERO = 00-20 "SOURCE OF CONSTANT ZERO FOR G SPEC SELECTION"

...S. G(G.ZERO), AUX0(1)=GSPEC, FSEL0=FIST, "CLEAR FIST"

FSEL1(1)=F.ZERO, AUX1=A, "CLEAR A OF R31"

3. SELECTING ADDRESS REFERENCES:

... NEXT.INSTR = N. + 1 "N. == CURRENT NANO-LOCATION-COUNTER"

.... BRANCH (NEXT.INSTR)

"OR OPTIONALLY"

.... BRANCH (N. + 1) "SETS KN = N. + 1"

6.3.2.4 6 BIT DATA TRANSFER COMMANDS

SPECIFICATION OF 6 BIT DATA TRANSFERS BETWEEN AUX FIELDS AND F REGISTERS, AND DIRECT F REGISTER MODIFICATION, REQUIRES THE USE OF A MINIMUM OF 3 T-VECTOR NANOPRIMITIVE FIELDS. IN ADDITION, ONLY THE A, B, AND KA FIELDS ARE UNIFORMLY ACCESSIBLE AS SOURCES AND DESTINATIONS IN ALL 3-6 BIT CONTROL GROUPS.

THE NANOPROGRAMMER HAS THE RESPONSIBILITY OF KNOWING HOW MANY OCCURRENCES OF EACH AUX EXIST IN A T-VECTOR, BUT THE ASSEMBLER CAN DETERMINE WHICH CONTROL GROUP TO USE FOR EACH 6 BIT DATA OPERATION. THIS AUTOMATIC SELECTION IS IN EFFECT WHEN USING THE 6 BIT DATA TRANSFER COMMANDS. THERE ARE FOUR FORMATS AVAILABLE FOR STATING THESE COMMANDS, AS FOLLOWS.

1. (SOURCE AUX NAME)->(EXPRESSION)
2. (SOURCE AUX NAME)->(EXPRESSION)->(DESTINATION AUX NAME)
3. (SOURCE AUX NAME)->(DESTINATION AUX NAME)
4. (EXPRESSION)->(DESTINATION AUX NAME)

[VERSION 1, LEVEL 2, IS NOT COMPATIBLE WITH THE ABOVE DEFINITION. TRANSFER COMMAND DECLARATIONS MUST BE STATED, BY THE NANOPROGRAMMER, IN THE BEST FIT, LEFT TO RIGHT, ORDER, FOR CORRECT PLACEMENT INTO CONTROL GROUPS.]

[VERSION 1, LEVEL 2, (EXPRESSION) MAY NOT BE USED, AS ONLY A SINGLE VARIABLE SYMBOL NAME WILL BE RECOGNIZED IN THESE FIELDS. SEE EXAMPLE BELOW.]

(EXPRESSION) IS EQUIVALENT TO ALL LEGAL ARITHMETIC EXPRESSIONS ALLOWED IN THE EXPLICIT F-SELECT FIELD COMMANDS, AS DESCRIBED IN SECTION 6.3.2.2. BOTH THE (SOURCE AUX NAME) AND (DESTINATION AUX NAME) COMPONENTS MUST BE ONE OF THOSE LISTED IN THE TABLES BELOW. FORMAT 1 SPECIFIES AN AUX FIELD TO F REGISTER TRANSFER, OR DIRECT F REGISTER MODIFICATION (IE. INCREMENT F). FORMAT 4 SPECIFIES AN F REGISTER TO AUX FIELD TRANSFER. FORMAT 2 SPECIFIES A REPLACE-MENT OPERATION, WHERE THE AUX TRANSFER PARALLELS THE F TRANSFER, USUALLY EXCHANGING THE TWO FIELDS. IT MAY ALSO INDICATE A PASS OPERATION, WHEN THE VALUE OF (EXPRESSION) IS FIPH, WHERE THE SOURCE AUX FIELD IS TRANSFERED DIRECTLY TO THE DESTINATION AUX FIELD. FINALLY, FORMAT 3 IS EQUIVALENT TO FORMAT 2, SPECIFICALLY INVOKING THE USE OF FIPH TO PASS ONE AUX FIELD TO ANOTHER.

THE FOLLOWING TABLE LISTS ALL SOURCE AUX NAMES AND THEIR ACTUAL CONTROL GROUP OCCURRENCES:

SOURCE AUX NAME	VALUE	GROUP NUMBERS		
A	0	0	1	2
B	1	0	1	2
C	2		1	
KA	3	0	1	2
KB	4	0		2
KX	2			2
KT	4		1	
INCF	6		1	2
DECF	7		1	2
IO ID	7	0		
ALUF	6	0		
G	5	0	1	2

THE FOLLOWING TABLE LISTS ALL DESTINATION AUX NAMES AND THEIR ACTUAL CONTROL GROUP OCCURRENCES:

DESTINATION AUX NAME	VALUE	GROUP NUMBERS		
A	0	0	1	2
B	1	0	1	2
C	2	0	1	2
KA	3	0	1	2
KB	4	0	1	2
KS	7	0		
KX	5		1	
KT	6			2
KSHC	5	0		2
KALC	6	0	1	
KSHA	7		1	2

EXAMPLES OF TRANSFER COMMANDS:

1. TRANSFER THE A FIELD OF R31 TO FMOD.
A->FMOD "EQUIVALENT TO: AUX0(1)=A, FSEL0=FMOD"
2. EXCHANGE THE CONTENTS OF THE B FIELD WITH FIST.

- B->FIST->B "EQUIVALENT TO: AUX0(1)=B, FSEL0(1)=FIST"
3. TRANSFER KT TO G2, WHILE TRANSFERRING G2 TO KSHA.
KT->G2->KSHA "EQUIVALENT TO: AUX1(1)=KT, FSEL1=G2, "
" FSEL2(1)=G2, AUX3=KSHA"
4. TRANSFER KX DIRECTLY TO KALC.
KX->KALC "SAME AS KX->FIPH->KALC"
5. TRANSFER AN F REGISTER ONE GREATER THAN F.WORK TO KX.
F.WORK+1->KX "EQUIVALENT TO: FSEL0(1)=F.WORK+1, AUX0=KX."
- I. VERSION 1, LEVEL 2 RESTRICTION, (EXPRESSIONS) MAY NOT BE USED IN 6 BIT TRANSFER COMMANDS. ONLY SINGLE SYMBOLIC NAMES MAY BE REFERENCED. IN ORDER TO ACCOMPLISH THAT SHOWN IN EXAMPLE 5, A TEMPORARY NAME MUST BE USED TO HOLD THE F REGISTER VALUE. (IE. TEMP=F.WORK+1, TEMP->KX). J

6.3.3 CONTROL STATEMENTS

CONTROL STATEMENTS CONSIST OF ONE OR MORE CONTROL FIELDS. THOUGH ALL CONTROL FIELD OPERATORS WILL BE RECOGNIZED ON OTHER STATEMENT TYPES, THE FOLLOWING CAUTION SHOULD BE OBSERVED. ALL CONTROL FIELDS TERMINATE THE NANOWORD CURRENTLY BEING ASSEMBLED. IN ADDITION, STATEMENTS CONTAINING SOME OF THE CONTROL FIELDS WILL NOT BE LISTED. TWO CLASSES OF CONTROL STATEMENTS ARE DEFINED FOR VERSION 1: ASSEMBLY CONTROL, AND LISTING CONTROL.

6.3.3.1 ASSEMBLY CONTROL STATEMENTS

ONLY ONE ASSEMBLY CONTROL STATEMENT IS DEFINED FOR VERSION 1. THE "END" CONTROL OPERATOR INDICATES THE END OF THE LAST SOURCE STATEMENT OF THE CURRENT ASSEMBLY. IT SHOULD BE PLACED BY ITSELF ON A SOURCE STATEMENT. IT WILL ALWAYS BE LISTED, REGARDLESS OF LISTING CONTROLS SPECIFIED.

6.3.3.2 LISTING CONTROL STATEMENTS

THERE ARE 4 LISTING CONTROL STATEMENTS DEFINED FOR VERSION 1. EACH TERMINATES THE CURRENT NANOWORD. SINCE SOME ARE NOT LISTED THEY SHOULD BE PLACED ON THEIR OWN SOURCE STATEMENTS.

LIST OFF NO PARAMETERS. REMAINDER OF STATEMENT IS PROCESSED. THE STATEMENT CONTAINING "LIST OFF" WILL NOT BE LISTED, AS WELL AS ALL FOLLOWING STATEMENTS UNTIL ONE WITH A "LIST ON" CONTROL COMMAND

IS ENCOUNTERED. LINES IN ERROR WILL BE LISTED UNCONDITIONALLY.

LIST ON NO PARAMETERS. REMAINDER OF STATEMENT IS PROCESSED. REVERSES THE EFFECT OF A PREVIOUS "LIST OFF" COMMAND.

". " (SINGLE PERIOD) NO PARAMETERS. REMAINDER OF STATEMENT IS PROCESSED. ". " SIMPLY INDICATES THE END OF THE CURRENT NANOWORD. IT IS USED OPTIONALLY TO TRIGGER NORMAL ASSEMBLY GENERATED END OF WORD LISTING INFORMATION. THIS PERMITS INSERTION OF EXTRA LINES AND COMMENTS AHEAD OF THE NEXT WORD DEFINITION. WHEN NOT TERMINATED BY COMMAND, THE LABEL STATEMENT OF THE NEXT NANOWORD TRIGGERS THE GENERATED LISTING INFORMATION.

EJECT NO PARAMETERS. REMAINDER OF STATEMENT IS IGNORED. "EJECT" IS NEVER LISTED BUT WILL CAUSE AN EJECT TO TOP OF NEXT PAGE, WHILE "LIST ON" IS IN EFFECT.

6.4 OPERATION [VERSION 1, LEVEL 2, UNDER NCS ONLY]

6.4.1 INVOCATION

THE NANO-ASSEMBLER IS INITIATED BY COMMAND AT THE SYSTEM CONSOLE. ENTERING THE NAME "NA," WILL BEGIN ASSEMBLER EXECUTION. THE REQUEST FOR INPUT FILE NAME WILL BE DISPLAYED AS "INPT=". RESPOND WITH THE DISK DATA FILE NAME, FOLLOWED BY A COMMA. THE REQUEST FOR BINARY OUTPUT FILE NAME WILL BE DISPLAYED NEXT AS "BIN=". RESPOND WITH EITHER THE FILE NAME TO BE USED, FOLLOWED BY A PERIOD, OR JUST A PERIOD TO INDICATE THAT NO OUTPUT FILE IS DESIRED. THE ESCAPE KEY (ESC) MAY BE USED TO CANCEL A PARTIALLY ENTERED FILE NAME, IN ORDER TO CORRECT KEYING MISTAKES.

IF ANY ERRORS ARE DETECTED DURING THE ASSEMBLY THE MESSAGE "ASSEM. ERRORS" WILL BE DISPLAYED AT THE END OF PROCESSING.

6.4.2 ERROR FLAGS

ERRORS ARE INDICATED THROUGH USE OF SINGLE CHARACTER CODES PLACED ON THE LINE IMMEDIATELY FOLLOWING THE LINE IN ERROR. EACH CODE WILL APPEAR DIRECTLY UNDER THE SYMBOL OR CHARACTER IN ERROR. THERE ARE THREE CLASSES OF ERROR DETECTION: LEXICAL, SYNTACTIC, AND GENERAL.

LEXICAL ERRORS ARE DETECTED DURING INITIAL SCAN OVER THE SOURCE STATEMENTS. ANY ILLEGAL CHARACTERS OR UNRECOGNIZED CHARACTER SEQUENCES ARE FLAGGED WITH A DIGIT AS FOLLOWS:

CODE	DESCRIPTION
------	-------------

0	ILLEGAL FIRST CHARACTER IN A FIELD.	
1	ILLEGAL CHARACTER WITHIN A SYMBOLIC NAME.	
2	ILLEGAL CHARACTER WITHIN A NUMERIC STRING.	
3	ILLEGAL OCTAL NUMBER.	
4	INTERNAL ERROR IN LEXICAL ANALYZER.	***
5	EXPECTED OPERATOR OR FIELD DELIMITER MISSING.	
6	SYNTAX TABLE FULL (EXPRESSION OVERLY COMPLEX)	***

*** REPORT THESE ERRORS TO: THE SYSTEMS SOFTWARE DIVISION,
NANODATA CORPORATION, 2457 WEHRLE DRIVE, WILLIAMSVILLE, NEW YORK 14221.
PLEASE PROVIDE AS MUCH SUPPORTING MATERIAL AS POSSIBLE (WITHIN REASON),
LISTINGS, DECKS, DUMPS, ETC.

LEXICAL ERRORS CAUSE A SKIP TO THE NEXT COMMA, OR END OF STATEMENT, WHICHEVER OCCURS FIRST. THIS SKIP IS INDICATED BY A STRING OF HYPHENS FROM THE ERROR CODE TO THE END OF THE SKIP, ON THE ERROR FLAG LINE.

SYNTACTIC ERRORS ARE DETECTED WHILE THE ASSEMBLER IS ATTEMPTING TO CLASSIFY EACH FIELD (NANOPRIMITIVE, LABEL, ARITHMETIC EXPRESSION, CONTROL, ETC.). EACH ERROR RECEIVES A LETTER CODE AS FOLLOWS:

CODE	DESCRIPTION
A	IMPROPER FIRST ELEMENT IN A FIELD.
B	FIRST ELEMENT IMPROPERLY TERMINATED.
C	IMPROPER FIRST ELEMENT IN A NANOPRIMITIVE SECONDARY FIELD.
D	IMPROPER FIRST ELEMENT IN AN ARITHMETIC EXPRESSION.
E	ILLEGAL ELEMENT, OR IMPROPER TERMINATION, IN A NANOPRIMITIVE SECONDARY FIELD.
F	IMPROPER ELEMENT IN SECONDARY FIELD.
G	SECONDARY FIELD NOT FOLLOWED BY LEGAL DELIMITER.
H	IMPROPER ELEMENT WITHIN AN ARITHMETIC EXPRESSION.
I	IMPROPER ARITHMETIC OPERATOR WITHIN AN EXPRESSION.
J	ILLEGAL COMPONENT IN A 6 BIT DATA TRANSFER COMMAND.
K	" " " " " " " "
L	" " " " " " " "
M	IMPROPERLY FORMATTED, THREE ELEMENT, 6 BIT DATA TRANSFER COMMAND.

GENERAL ERROR FLAGS ARE SET BY VARIOUS STATEMENT CLASS, AND FIELD TYPE, PROCESSES. THESE ARE LISTED BELOW:

CODE	DESCRIPTION
?	UNRECOGNIZED NANOPRIMITIVE COMMAND, CONTROL OPERATOR, OR PSEUDO COMMAND OPERATOR.
=	WARNING; THAT A REDEFINITION OF ONE OF THE PREDEFINED REFERENCE VARIABLES HAS OCCURRED.
\$	INTERNAL ASSEMBLER CONTROL ERROR. *** (ABOVE)
M	MULTIPLY DEFINED SYMBOLIC LABEL NAME, OR MULTIPLY DECLARED NANOPRIMITIVE FIELD.
T	T-VECTOR NANOPRIMITIVE DECLARED WITHIN K-VECTOR STATEMENT RANGE.

```

N      NANOPRIMITIVE DECLARED OUTSIDE THE RANGE OF A NANOWORD.
P      6 BIT DATA TRANSFER COMMAND WILL NOT FIT IN THE CURRENT T-VECTOR.
      ALL ACCEPTABLE GROUPS IN USE.
X      UNDEFINED VARIABLE NAME IN ARITHMETIC EXPRESSION.
U      UNNECESSARY PRIMARY OR SECONDARY NANOPRIMITIVE COMMAND SPECIFICATION.
I      INCOMPLETE NANOPRIMITIVE SPECIFICATION.  A REQUIRED PRIMARY OR
      SECONDARY FIELD DECLARATION IS MISSING.

```

GENERAL: LISTABLE ERROR MESSAGES:

[illegible]

CONSOLE ERROR MESSAGES:

BAD NAME	- ILLEGAL DISK FILE NAME ENTERED.
FILE NOT FOUND	- TRY ANOTHER NAME, DISK OR VOLUME.
KALFUNCTION	- USUALLY ARE HARDWARE FAILURE DURING ASSEMBLY.
SYMBOL TABLE OVERFLOW	- ASSEMBLY REQUIRES MORE MEMORY SPACE.
ASSEM. ERRORS	- ONE OR MORE ERRORS DURING ASSEMBLY.

2 NANOPROGRAMMING EXAMPLES

2.1 BASIC MODEL NANOPROGRAMS

ONE OF THE MOST BASIC NANOPROGRAMS THAT CAN BE ACTUALLY USED IS ILLUSTRATED BELOW. THIS NANOPROGRAM IMPLEMENTS A WAIT MICROINSTRUCTION THAT CYCLES UNTIL AN INTERRUPT OCCURS.

WAIT:

```
.... LEGAL MICRO OP ENTRY, ALLOW NANO INTERRUPT, ALLOW MICRO INTERRUPT
X...
.X.. READ NS
..X. GATE NS
...X
```

THE K-VECTOR COMMANDS SPECIFY THAT THIS IS A LEGAL MICROINSTRUCTION AND THAT ALL INTERRUPTS ARE ALLOWED FOLLOWING EXECUTION OF THIS NANOWORD.

IN THE SECOND T-STEP, NANOSTORE IS READ USING THE ADDRESS THAT REMAINS IN THE NANOPROGRAM COUNTER (NPC) FOLLOWING THE BEGINNING OF EXECUTION OF THE NANOWORD. IN THE THIRD T-STEP, THE NANOSTORE WORD READ IS GATED INTO THE CONTROL MATRIX. IT WILL BEGIN EXECUTION FOLLOWING THE THIRD T-STEP (THE LAST T-STEP IS UNUSED.)

SINCE THE SUCCESSIVE NANOSTORE WORDS READ ARE THE SAME AS THE FIRST ONE TRIGGERED BY EXECUTION OF THE WAIT MICROINSTRUCTION, THIS NANOPROGRAM CYCLES REPEATEDLY UNTIL AN INTERRUPT OCCURS.

WHEN AN INTERRUPT OCCURS (AND IS ACCEPTED BY THE MACHINE), THE ABOVE SEQUENCE IS SUSPENDED. THE NEXT NANOWORD READ AND PLACED INTO EXECUTION IS THE ONE ADDRESSED BY THE PARTICULAR INTERRUPT THAT HAS OCCURRED.

THE FIRST T-STEP IS NOT USED FOR THE READ NS NANOPRIMITIVE, IN THE EXAMPLE, SINCE INTERRUPTS ARE ALLOWED FOLLOWING THE EXECUTION OF THIS NANOWORD. INTERRUPT ADDRESS SELECTION REQUIRES AT LEAST ONE T-PERIOD, AND CANNOT BEGIN UNTIL AT LEAST ONE OF THE ALLOW INTERRUPT K-VECTOR NANOPRIMITIVES IS RECOGNIZED. THE EARLIEST INTERRUPT ALLOW RECOGNITION TIME IS AT THE LEADING EDGE OF T1. SINCE READ NS ADDRESS SELECTION IS ALSO A LEADING EDGE FUNCTION IT MAY NOT BE EXECUTED IN T1, UNLESS IT IS SELECTING A NANOBANCH (KN) ADDRESS.

IN THE EXAMPLE ABOVE, NO ACCESS IS MADE TO CONTROL STORE SINCE THE SAME NANOWORD IS REPEATEDLY READ. CONSEQUENTLY, NOTHING HAS TO BE DONE TO CHANGE THE

MICROPROGRAM COUNTER. ITS CURRENT VALUE STILL INDICATES THE CONTROL STORE WAIT INSTRUCTION THAT INITIATED THE NANOPROGRAM. IN THE MORE TYPICAL SITUATION, SUCCESSIVE CONTROL STORE INSTRUCTIONS WOULD BE READ AND USED TO ADDRESS DIFFERENT NANOPROGRAMS THAT IMPLEMENT THE SUCCESSIVE INSTRUCTIONS. THE NEXT EXAMPLE SHOWS THE IMPLEMENTATION OF A NO-OPERATION MICROINSTRUCTION THAT INCLUDES READING THE NEXT MICROINSTRUCTION FROM CONTROL STORE AND UPDATING THE MICROPROGRAM COUNTER.

NOP:

```
.... LEGAL MICRO OP ENTRY, ALLOW NANO INTERRUPT, ALLOW MICRO INTERRUPT
X... READ CS(MPC+1)
..X.. LOAD NPC(CS), MPC PLUS 1  "INC MPC +1; THIS IS A COMMENT"
...X. READ NS
....X GATE NS, LOAD R31
```

IN THIS NANOPROGRAM, THE CONVENTION HAS BEEN ESTABLISHED THAT AT NANOPROGRAM ENTRY, THE MPC (ONE OF 4 AVAILABLE LOCAL STORE REGISTERS) POINTS TO THE CURRENTLY EXECUTING MICROINSTRUCTION IN CONTROL STORE. THE ACTUAL MPC IN USE IS DETERMINED BY THE CONTENTS OF FMPC WHEN THE NANOPROGRAM IS EXECUTED. IT IS FURTHER DEFINED THAT THE NOP MICROINSTRUCTION IS ONE WORD LONG. (THIS IS A REASONABLE ASSUMPTION SINCE NO PARAMETERS ARE NEEDED IN A NOP!) THUS DURING T1, CONTROL STORE IS READ AT THE ADDRESS ONE GREATER THAN THE CURRENT CONTENTS OF THE MICROPROGRAM COUNTER. THUS THE NEXT MICROINSTRUCTION IN SEQUENCE IS BEING READ.

DURING T2, THE MICROPROGRAM COUNTER IS UPDATED. THIS ASSURES THAT IT WILL POINT TO THE NEXT MICROINSTRUCTION WHEN THAT MICROINSTRUCTION BEGINS EXECUTION. SIMULTANEOUSLY, THIS NEW MICROINSTRUCTION FROM CONTROL STORE (ALONG WITH THE CONTENTS OF THE PAGE REGISTER IN FIDX) IS USED TO LOAD THE NANOPROGRAM COUNTER (NPC).

FROM THIS POINT ON, THE NANOPROGRAM IS SIMILAR TO THE FIRST EXAMPLE. THE NANOSTORE WORD THAT BEGINS THE IMPLEMENTATION OF THE NEXT MICROINSTRUCTION IS READ AND GATED INTO THE CONTROL MATRIX. SIMULTANEOUS WITH THIS LAST ACTION, THE A AND B PARAMETERS OF THE NEW MICROINSTRUCTION ARE LOADED INTO R31. THE NEW NANOPROGRAM THEN BEGINS EXECUTION.

THIS EXAMPLE IS SHOWN AS A MODEL SINCE IT CONSISTS OF A BASIC SET OF NANO-PRIMITIVES THAT WILL BE COMMON TO MANY NANOPROGRAMS. IT FORMS A BASIS FOR ALL ONE WORD NANOPROGRAMS THAT IMPLEMENT ONE-WORD MICROINSTRUCTIONS UNDER THE NANOPROGRAMMING CONVENTIONS MENTIONED. THE NEXT EXAMPLE WILL ILLUSTRATE ANOTHER OF THE MANY POSSIBLE SETS OF CONVENTIONS THAT MAY BE SELECTED.

ONE POSSIBLE NANOPROGRAMMING CONVENTION WHICH CAN BE ESTABLISHED IS THAT EACH NANOPROGRAM WOULD EXPECT THE NEXT SEQUENTIAL WORD FROM CONTROL STORE TO BE AVAILABLE AT THE BEGINNING OF ITS EXECUTION (ON THE COD BUS) AND THAT EACH NANOPROGRAM WOULD BE RESPONSIBLE FOR MAINTAINING THIS PREFETCH CONVENTION FOR THE NEXT NANOPROGRAM. THE FOLLOWING EXAMPLE SHOWS A POSSIBLE FORM OF THE NOP WHEN PROGRAMMED ACCORDING TO THIS NEW CONVENTION.

NOPLA:

```
....  LEGAL MICRO OP ENTRY, ALLOW INTS
X...  LOAD NPC(CS)
.X...  READ NS, READ CS(MPC+2), INC MPC
..X... GATE NS, LOAD R31
```

SINCE THE NEXT WORD FROM CONTROL STORE IS ALREADY AVAILABLE, WE CAN IMMEDIATELY LOAD THE NANOPROGRAM COUNTER IN T1. AT THE SAME TIME, WE START THE OPERATIONS NECESSARY TO SET UP A SIMILAR SITUATION AT THE END OF THE NANOPROGRAM BY READING CONTROL STORE. THE ADDRESS USED IS MPC+2 SINCE WE NEED TO READ THE WORD AHEAD BY TWO FROM THE CURRENTLY EXECUTING MICROINSTRUCTION.

THE REMAINDER OF THE EXAMPLE PROCEEDS AS BEFORE. AT THE END OF T3, WE ARE READY TO EXECUTE THE NEXT NANOPROGRAM AND THE NEXT SEQUENTIAL WORD FROM CONTROL STORE IS AVAILABLE ON THE COD BUS. AS BEFORE, THE MICROPROGRAM COUNTER POINTS TO THE CURRENTLY EXECUTING MICROINSTRUCTION.

THIS EXAMPLE HAS BEEN GIVEN FOR ILLUSTRATIVE PURPOSES ONLY; THE ADVANTAGES OF A CONTROL STORE PREFETCH ARE NOT EXPLORED HERE. IT DOES ILLUSTRATE, HOWEVER, SOME OF THE POSSIBLE FREEDOM AVAILABLE TO THE NANOPROGRAMMER IN SELECTING THE CONVENTIONS THAT BEST SUIT HIS PURPOSES.

NOW AN EXAMPLE WILL BE GIVEN SHOWING A VERY SIMPLE NANOPROGRAM WHICH ACTUALLY USES THE MICROINSTRUCTION PARAMETERS. THIS IS AN IMPLEMENTATION OF THE MOV A,B MICROINSTRUCTION THAT CAUSES THE CONTENTS OF LOCAL STORE REGISTER B TO BE MOVED TO LOCAL STORE REGISTER A, WHERE A AND B ARE THE PARAMETERS IN THE MOV MICROINSTRUCTION.

MOV:

```
....  LEGAL MICRO OP, ALLOW INTS
....  KSHA=0, KSHC=0
X...  READ CS(MPC+1), A->FSOD, B->FSID
.X...  LOAD NPC(CS), MPC PLUS 1
..X... READ NS, GATE SH
...X... GATE NS, LOAD R31
```

THIS NANOPROGRAM USES THE SHIFTER AS A PATH TO ACCOMPLISH THE MOVE. THUS THE MAIN PART OF THE PROGRAM IS THE SETTING OF THE SHIFTER BUS CONTROLS FSID AND FSOD, FOLLOWED BY THE GATING OF THE SHIFTER TO ACTUALLY CAUSE THE TRANSFER. NOTE THAT THE CONTROLS ARE SET UP IN T1 BUT THE GATE OPERATION IS DEFERRED UNTIL T3. THIS IS NECESSARY IN ORDER FOR THE SHIFTER OUTPUT TO BE STABLE.

SINCE THE SHIFTER IS USED MERELY AS A PATH TO ACCOMPLISH THE MOVE, NO SHIFTING IS REQUIRED. THE SHIFT AMOUNT IS EXPLICITLY SET TO ZERO WITH KSHA=0. THIS IS DONE TO MAKE THE EXAMPLE CLEAR; IT IS UNNECESSARY, AS IS KSHC=0, SINCE THE DEFAULT VALUE FOR THESE FIELDS IS ZERO WHEN THEY ARE NOT MENTIONED.

NOTE THAT ALL OF THE NANOPRIMITIVES OF THE NOP NANOPROGRAM ARE INCLUDED IN THE EXAMPLE MOV NANOPROGRAM. THIS IS DONE SINCE SOME SCHEME FOR FETCHING AND SEQUENCING OF MICROINSTRUCTIONS IS NECESSARY. THE SCHEME ILLUSTRATED IN THE NOP NANOPROGRAM IS ONE SUCH SCHEME AND IS SATISFACTORY FOR THE INSTRUCTION PRESENTED. THE PREFETCH SCHEME WOULD WORK AS WELL AS IS SHOWN BELOW. - -

MOVLA:

```
....  LEGAL MICRO, ALLOW INTS
....  KSHA=0, KSHC=0
X...  LOAD NPC(CS), A->FSOD, B->FSID
.X...  READ NS, READ CS (MPC+2), MPC PLUS 1
..X... GATE NS, GATE SH, LOAD R31
```

IN FACT, THIS ALSO ILLUSTRATES A POSSIBLE ADVANTAGE OF THE PREFETCH SINCE THE NANOPROGRAM IS SHORTER BY ONE PERIOD.

AT THIS POINT, IT IS USEFUL TO INTRODUCE A SHORTHAND NOTATION TO SIMPLIFY WRITING SETS OF NANOPROGRAMS. BY PREDEFINING A SET OF NANOPRIMITIVES THAT APPEAR FREQUENTLY, THEY MAY BE INVOKED BY NAME WHEN NEEDED. THIS IS ILLUSTRATED IN THE NEXT SECTION.

7.2 USE OF PREDEFINED NANOWORDS

[NOTE: IN CURRENT VERSIONS OF THE NANO-ASSEMBLER NO SUPPORT IS PROVIDED]
[FOR FETCHING PREDEFINED NANOWORDS. IT IS LIKELY THAT THIS SUPPORT WILL BE]
[ADDED IN LATER VERSIONS. UNTIL THEN, PREDEFINED NANOWORDS ARE DESCRIBED IN]
[THIS DOCUMENT STRICTLY FOR ILLUSTRATIVE PURPOSES.]

IT IS USEFUL TO HAVE A WAY TO CONDENSE THE DESCRIPTION OF SETS OF NANOPROGRAMS HAVING FREQUENT REPETITION OF THE SAME NANOPRIMITIVE SEQUENCES. THIS IS DONE BY REFERENCE TO PREDEFINED NANOWORDS. ANY NANOWORD PREVIOUSLY DEFINED AND LABELED MAY BE INVOKED TO CAUSE ALL OF THE PREDEFINED BITS TO BE SET IN THE NANOWORD IN WHICH THE LABEL OF THE PREDEFINED NANOWORD IS MENTIONED.

FOR EXAMPLE, THE MOV NANOPROGRAMS OF THE PREVIOUS SECTION MAY BE WRITTEN:

```
MOV:
.... NOP, KSHA=0, KSHC=0
X... A->FSOD, B->FSID
..X. GATE SH
```

HERE THE FETCH SEQUENCE IS INVOKED BY REFERENCE TO NOP, PROVIDED THE NOP NANOWORD HAS BEEN DEFINED AS SHOWN IN THE PREVIOUS SECTION. ALL OF THE BITS ON IN THAT NANOWORD WILL BE SET ON IN THIS NANOWORD.

SIMILARLY, THE PREFETCH SEQUENCE CAN BE INVOKED INSTEAD BY CHANGING NOP TO NOPLA.

```
NOVLA:
.... NOPLA, KSHA=0, KSHC=0
X... A->FSOD, B->FSID
..X. GATE SH
```

THIS NOTATION WILL BE USED EXTENSIVELY IN THE EXAMPLES THAT FOLLOW. IT MAKES THE EXAMPLES SHORTER AND ALLOWS ONE TO CONCENTRATE ON THE PARTS OF EACH NANOPROGRAM THAT ARE NOVEL TO EACH NEW EXAMPLE. CONSIDERABLE CARE MUST BE EXERTED, HOWEVER, TO AVOID COMBINATIONS THAT ARE NOT CONSISTENT.

THE SIMPLE FETCH PROCEDURE, DESCRIBED IN THE BEGINNING OF THIS CHAPTER, WILL BE USED IN MOST OF THE EXAMPLES. IT IS FORMALLY CODED AS FOLLOWS:

FETCH:

```
....  LEGAL MICRO OP ENTRY, ALLOW NANO INTERRUPT, ALLOW MICRO INTERRUPT
X...  READ CS (MPC+1)
..X..  LOAD NPC (CS)
...X.  READ NS, MPC PLUS 1
...X  GATE NS, LOAD R31
```

THE INC MPC NANOPRIMITIVE (MPC PLUS 1) HAS BEEN PLACED IN T3 SO THAT THE PREVIOUS CONTENTS OF MPC MAY BE USED DURING T1, T2 OR T3 IN NANOPROGRAMS THAT USE FETCH.

TO ILLUSTRATE THE USE OF THIS PREDEFINED NANOWORD, THE FOLLOWING NANOPROGRAM IMPLEMENTS A LD A,B MICROINSTRUCTION THAT LOADS FROM CONTROL STORE, INTO LOCAL STORE REGISTER A, THE WORD ADDRESSED BY THE CONTENTS OF LOCAL STORE REGISTER B.

LD:

```
....  FETCH
X...  A->FCOD, B->FCIA
..X..  READ CS (CIA)
...X  GATE CS
```

IN THIS EXAMPLE, THE ESSENTIAL BUT PREVIOUSLY DISCUSSED PARTS OF THE NANOPROGRAM ARE CONDENSED INTO THE REFERENCE TO FETCH. THEN, THE PARTS OF THE NANOPROGRAM IMPORTANT TO MAKE THE LD WORK CAN MORE CLEARLY BE SEEN.

2.3 CONTROL STORE ACCESS NANOPROGRAMS

THE PREVIOUS EXAMPLE SHOWS HOW ONE MAY IMPLEMENT AN INSTRUCTION TO LOAD A WORD FROM CONTROL STORE INTO A LOCAL STORE REGISTER, USING AN ADDRESS ALREADY IN A LOCAL STORE REGISTER. THE EXAMPLES IN THIS SECTION WILL ILLUSTRATE OTHER WAYS IN WHICH CONTROL STORE MAY BE ADDRESSED AND ACCESSED.

THE ADDRESS OF THE DESIRED CONTROL STORE DATA LOCATION MAY ITSELF BE FETCHED FROM CONTROL STORE. TWO SIMPLE ACCESS TECHNIQUES CAN BE IMPLEMENTED. FIRST OF ALL, ONE COULD TREAT THE LOWER 64 (DECIMAL) LOCATIONS OF CONTROL STORE AS A SET OF DIRECTLY ACCESSED SPECIAL REGISTERS. SECOND, THE DATA ADDRESS COULD BE FOUND AT SOME CONTROL STORE LOCATION RELATIVE TO THE LOAD INSTRUCTION ITSELF. THE MOST ACCESSIBLE LOCATION, IN THIS CASE, WOULD BE THE NEXT SEQUENTIAL CONTROL STORE WORD.

THE FOLLOWING EXAMPLE ILLUSTRATES A MICROINSTRUCTION THAT READS THE CONTROL STORE LOCATION IDENTIFIED BY THE VALUE OF B, INTO LOCAL STORE REGISTER A.

```
LDCSR:
....  FETCH, KA=31.
X...  A->FCOD, KA->FCIA, FIPH->A
.X...
...X.  READ CS (CIA)
...X   GATE CS
```

THE LOAD-VIA-CONTROL-STORE-REGISTER INSTRUCTION USES THE PREDEFINED FETCH PROCEDURE TO ACCOMPLISH ITS NEXT MICROINSTRUCTION ACCESS. CARE MUST BE TAKEN NOT TO INTERFERE WITH THAT LOGIC, ESPECIALLY SINCE BOTH FETCH AND THE INSTRUCTION EXECUTION LOGIC ACCESS CONTROL STORE.

T1 OF OUR EXAMPLE SETS UP THE F-REGISTERS FOR THE DATA ACCESS. LOCAL STORE REGISTER A IS CONNECTED TO THE COD BUS, REGISTER 31 (DECIMAL) IS CONNECTED TO THE CIA BUS, AND THE ORIGINAL CONTENT OF THE A FIELD IS ZEROED LEAVING ONLY THE B FIELD REMAINING IN R31. FIPH IS A SOURCE OF ZEROES WHEN NO SOURCE AUX IS TRANSMITTED TO IT, WITHIN THE SAME T-STEP. NOTE THAT THE C FIELD OF R31 DOES NOT NEED TO BE CLEARED, AS THE CONVENTIONAL MICROINSTRUCTION INITIATION SEQUENCE CLEARS C WHILE LOADING A AND B (LOAD R31 NANOPRIMITIVE).

T1 AND T2 ARE USED BY THE FETCH SEQUENCE TO ACCESS CONTROL STORE. WE MAY USE CONTROL STORE FOR OUR DATA ACCESS BEGINNING IN T3. LOCAL STORE REGISTER 31 IS NOW USED FOR ADDRESSING CONTROL STORE. NOTE ALSO THAT WE ARE ACCESSING BOTH NANOSTORE AND CONTROL STORE SIMULTANEOUSLY DURING T3 AND T4. THE FETCH

SEQUENCE IS READING THE NEXT NANOPROGRAM WHILE THE LDCSR LOGIC IS READING THE DATA WORD.

TO BE ABLE TO ADDRESS ANY CONTROL STORE LOCATION THE LOAD-VIA-NEXT-WORD INSTRUCTION MAY BE DEFINED. IN THIS CASE THE MICROINSTRUCTION MAY BE CONSIDERED AS BEING TWO WORDS IN LENGTH, WHERE THE SECOND WORD CONTAINS THE ABSOLUTE, 16 BIT CONTROL STORE ADDRESS. THE DATA IS LOADED INTO THE LOCAL STORE REGISTER IDENTIFIED BY THE A FIELD OF THE FIRST WORD OF THE INSTRUCTION.

THERE ARE MANY WAYS OF ORGANIZING THIS NANOPROGRAM. IN ORDER TO USE THE FETCH PREDEFINITION WE MUST DEFINE AN ADDITIONAL PROGRAMMING CONVENTION. IF WE ENFORCE THE RULE THAT ANY NANOPROGRAM THAT USES FCOD MUST RESTORE IT TO THE VALUE 31 (DECIMAL) PRIOR TO COMPLETION, WE MAY THEN EXECUTE A GATE CS NANO-PRIMITIVE IN T1 WITH THE KNOWLEDGE THAT THE COD BUS IS ALREADY CONNECTED TO R31.

```
LDNN:  "WITH FCOD=31 CONVENTION"
....  FETCH, KA=31.
S...  GATE CS, A->FCOD, KA->FCIR
.S...  READ CS (MPC+2), MPC PLUS 1
..X... READ CS (CIR)
...X   GATE CS, KA->FCOD
```

THE FETCH SEQUENCE NORMALLY INITIATES THE READ OF CS LOCATION MPC+1 IN T1, THEN LOADS NPC WITH THE RESULTS OF THE READ IN T2. BY STRETCHING T1 THE READ CS (MPC+1) IS COMPLETED WITHIN T1, AND OUR GATE CS WILL PLACE THE CONTENT OF MPC+1 INTO LOCAL STORE REGISTER 31 (BY OUR FCOD CONVENTION). THUS, BY THE END OF T1 OUR DATA ADDRESS RESIDES IN R31. SINCE WE HAVE CHANGED R31 BY THE END OF T1 WE MUST EXTRACT ALL NECESSARY A AND B FIELD INFORMATION DURING T1. B IS NOT USED, BUT A MUST BE TRANSFERED TO FCOD (WE HAVE NOW ALTERED FCOD).

THE FETCH SEQUENCE EXPECTS THE NEXT INSTRUCTION TO BE AVAILABLE ON THE COD BUS BY THE END OF T2. WE CAN PROVIDE THIS BY STRETCHING T2, AND EXECUTING A READ CS(MPC+2) AT THE BEGINNING OF THAT T-STEP. T3 AND T4 NOW CAN FUNCTION THE SAME AS IN THE LDCSR INSTRUCTION ABOVE. THE ONLY DIFFERENCE WILL BE OUR RESTORATION OF FCOD TO THE VALUE 31, AS REQUIRED BY OUR NEW CONVENTION.

IT IS POSSIBLE TO DEFINE AN LDNN INSTRUCTION WITHOUT THE ABOVE CONVENTION, BUT THE FETCH DEFINITION MAY NOT BE APPLIED. WE ALSO WILL REQUIRE TWO ADDITIONAL T-PERIODS AS FOLLOWS.

LDNW: "WITH NO CONVENTIONS"

```
*      INSTRUCTION EXECUTION CODE "" INSTRUCTION FETCH CODE
....
S...  READ CS(MPC+1), R->FCOD      ""
.S...  READ CS(COD), GATE CS      ""
..S... "" READ CS(MPC+2), MPC PLUS 2, LOAD NPC(CS)
...S   "" READ NS, GATE NS, LOAD R31
```

IN T1 ABOVE WE READ THE WORD FROM MPC+1 ONTO THE COD BUS, BUT WE DO NOT GATE IT INTO LOCAL STORE. WE ALSO SET UP FCOD. THEN, IN T2, WE POINT THE CONTROL STORE ADDRESS SELECT DIRECTLY AT THE DATA ON THE COD BUS, AND USE THIS AS OUR DATA ADDRESS. T2 MUST COMPLETE THE DATA ACCESS IN ORDER TO FREE CONTROL STORE FOR A THIRD READ, WHICH MUST BEGIN IN T3 IF WE ARE GOING TO COMPLETE THIS NANO-PROGRAM WITHIN ONE NANOWORD.

NOTE THAT IN BOTH OF THE LDNW EXAMPLES WE PERFORM THREE READ CS OPERATIONS. ONE IS REQUIRED TO READ THE DATA ADDRESS, ANOTHER TO ACCESS THE DATA, AND THE THIRD TO FETCH THE NEXT INSTRUCTION. ONLY THE ORDER OF EVENTS HAS CHANGED BETWEEN THE TWO METHODS. A LDNWL (LOOK AHEAD) INSTRUCTION SEQUENCE COULD ALSO BE WRITTEN IN A SIMILAR FASHION, SAVING AT LEAST ONE T-PERIOD. WE LEAVE THE PROOF OF THIS AS AN EXERCISE FOR THE READER.

THE FOLLOWING EXAMPLE ILLUSTRATES THE WRITING OF CONTROL STORE. THE STCSR INSTRUCTION IS THE COUNTERPART OF LDCSR, AND WILL STORE THE LOCAL STORE REGISTER IDENTIFIED BY THE A FIELD, INTO THE CONTROL STORE LOCATION IDENTIFIED BY THE ABSOLUTE VALUE OF B.

STCSR:

```
....  FETCH, KA=31.
X...  R->FCID, KA->FCIA, FIPH->A
.X...
..X... WRITE CS (CIA)
...X
```

THAT'S ALL THERE IS TO THAT. FCID IS SET UP IN T1, AND THE WRITE CS IS PERFORMED DURING T3 AND T4. EVERYTHING ELSE IS THE SAME AS IN THE LDCSR INSTRUCTION, ABOVE.

2.4 MICROINSTRUCTION BRANCH NANOPROGRAMS

THERE ARE SEVERAL QM-1 HARDWARE FUNCTIONS THAT SIMPLIFY THE DEFINITIONS OF BRANCH, OR JUMP, INSTRUCTIONS. ONLY THOSE BASIC INSTRUCTIONS WILL BE EXPLORED IN THIS SECTION. MORE ADVANCED INSTRUCTION FORMATS (IE. BRANCH AND LINK) WILL BE DISCUSSED IN SECTION 7.6. THREE BRANCHING EXAMPLES ARE SHOWN BELOW. ALTHOUGH ALL OF THE NANOPRIMITIVES OF THE FETCH SEQUENCE ARE USED IN ONE OF THE EXAMPLES, THE FUNCTIONAL MEANINGS OF THE USE OF EACH PRIMITIVE DIFFERS EXTENSIVELY, THEREFORE NO PREDEFINITIONS ARE USED IN THAT EXAMPLE.

THE ABILITY TO INCREMENT THE MICRO-PROGRAM-COUNTER BY THE VALUES 1 OR 2, AND TO ADD THE TWO'S COMPLEMENT VALUES OF EITHER THE 6 BIT B FIELD OR 11 BIT A AND B FIELD CONCATENATION TO THE MPC, SIMPLIFY THE DEFINITIONS OF SOME FORMS OF BRANCH INSTRUCTIONS. THE FIRST EXAMPLE EXECUTES A BRANCH THAT IS RELATIVE TO THE ADDRESS OF THE BRANCH MICROINSTRUCTION ITSELF. BPREL (BRANCH PROGRAM-COUNTER RELATIVE) ALLOWS A FORWARD BRANCH OF UP TO 1,023 (DECIMAL) LOCATIONS, AND A BACKWARD BRANCH OF UP TO 1,024 LOCATIONS. THE AB FIELD CONCATENATION HARDWARE PERMITS THE FOLLOWING ENCODING.

BPREL:

```
....  
X... "SETTLING TIME FOR COMPLETION OF MPC+AB CIRCUIT COMPUTATION."  
.S.. READ CS (MPC+AB "AB=4"), LOAD NPC (CS), MPC PLUS AB  
..X. READ NS  
...X GATE NS, LOAD R31
```

THE A AND B FIELDS OF R31 RECEIVE THE NEW MICROINSTRUCTION PARAMETERS AT THE START OF T1. THEREFORE SOME TIME MUST BE ALLOWED FOR THE NEW AB VALUE TO BE APPLIED TO THE MPC+AB ADDITION CIRCUIT. FOLLOWING T1 THE PROPER VALUE IS AVAILABLE FOR USE AS A CONTROL STORE ADDRESS AND FOR GATING INTO THE MPC REGISTER. T2 IS USED TO READ THE NEXT MICROINSTRUCTION, AND TO ROUTE ITS ADDRESS INTO MPC. NPC IS ALSO LOADED WITH THE NANOSTORE ADDRESS. T3 AND T4 THEN READ AND GATE THE NEW NANOPROGRAM, AS IN MOST PREVIOUS EXAMPLES.

FREQUENTLY THE BRANCH ADDRESS WILL BE FOUND IN A REGISTER. FOR EXAMPLE, ENTRY INTO A SUBROUTINE MAY LEAVE THE RETURN ADDRESS IN A REGISTER, REQUIRING A SIMPLE TRANSFER OF CONTROL TO THE INSTRUCTION AT THAT ADDRESS UPON SUBROUTINE EXIT. A BR (BRANCH-VIA-REGISTER) INSTRUCTION CAN BE DEFINED AS FOLLOWS.

```
OP
... KA=LS.MPC
    KA->FSOD, A->FSID, A->FCIA
    "ALLOW CIA AND SHIFTER SETTLING TIME."
    READ CS (CIA), LOAD NPC (CS), GATE SH
    READ NS, GATE NS, LOAD R31
```

IN THIS EXAMPLE THE MPC LOCAL STORE REGISTER HAS BEEN SYMBOLICALLY NAMED LS.MPC, AND IS DEFINED IN THE KA FIELD. SHOULD THE ACTUAL MPC BE A VARIABLE (REMEMBER, THERE ARE 4 POSSIBLE MPC REGISTERS) A CONVENTION MAY BE SPECIFIED WHEREBY THE CURRENT MPC REGISTER NUMBER COULD BE FOUND WITHIN A SPECIFIC G-REGISTER. IN THAT CASE IT COULD BE COPIED FROM THE G HOLDING REGISTER TO FSOD AS FOLLOWS.

```
S... G(G.MPC), G->FSOD, A->FSID, A->FCIA
```

THE SHIFTER IS BEING USED TO PASS THE ACTUAL ADDRESS FROM THE LOCAL STORE REGISTER, DEFINED BY A, INTO THE CURRENT MPC. CONTROL STORE ADDRESS SELECTION REQUIRES THAT THE ACTUAL ADDRESS BE AVAILABLE AT LEAST ONE T-PERIOD PRIOR TO THE ACTUAL READ CS OPERATION. THE ADDRESS WILL NOT BE STABLE, IN THE MPC, SOON ENOUGH TO BE REFERENCED BY A READ CS (MPC), ESPECIALLY IN A ONE NANOWORD PROGRAM. WE DO, THOUGH, HAVE THE ADDRESS AVAILABLE IN ITS ORIGINAL LOCAL STORE REGISTER, AND MAY REFERENCE THAT REGISTER FOR CS ADDRESSING VIA THE CIA BUS. THEREFORE A->FCIA IN T1 SETS UP THE APPROPRIATE CONNECTION SOON ENOUGH TO ALLOW OUR READ CS (CIA) IN T3.

THE COMPLETE CONTROL STORE ACCESS IS DONE IN A STRETCHED T3, ALLOWING NPC TO BE SET UP FOR USE BY THE READ NS IN T4. T4 IS ALSO STRETCHED, ALLOWING BOTH INITIATION AND COMPLETION OF THE NANOSTORE ACCESS.

ANOTHER FORM OF BRANCHING ALLOWS PROGRAMMED DECISION MAKING. THE NEXT EXAMPLE SHOWS A CONDITIONAL BRANCH INSTRUCTION. IN THIS CASE NO ADDRESS IS PASSED TO THE INSTRUCTION, AS IT REQUIRES NO OPERANDS. OUR INSTRUCTION WILL "SKIP" OVER THE NEXT SEQUENTIAL INSTRUCTION IF THE RESULT OF THE LAST ARITHMETIC OPERATION INSTRUCTION WAS POSITIVE, OR ZERO. THE SKIP-ON-PLUS INSTRUCTION WILL TEST THE SIGN INDICATOR OF F-REGISTER FIST. IF THE SIGN INDICATES A ZERO (POSITIVE) IT WILL SKIP THE NEXT INSTRUCTION. IF SIGN IS A ONE (NEGATIVE) IT WILL EXECUTE THE NEXT SEQUENTIAL INSTRUCTION.

SKIP ON PLUS:

```
.... KS = SIGN
*   FIRST READ THE NEXT SEQUENTIAL INSTRUCTION.
S... READ CS (MPC+1), LOAD NPC (CS), SKIP (NOT S)
*   SKIP OVER T2 IF (FIST .AND. KS) ARE NOT ZERO.
S... READ CS (MPC+2), LOAD NPC (CS), MPC PLUS 1
...S. READ NS,                MPC PLUS 1
...X GATE NS, LOAD R31
```

T1, ABOVE, READS THE NEXT SEQUENTIAL INSTRUCTION (MPC+1) AND PREPARES THE NPC IN CASE T2 IS SKIPPED. IF T2 IS NOT SKIPPED IT WILL READ THE MICROINSTRUCTION AT MPC+2, AND EXECUTE AN EXTRA INCREMENT MPC OPERATION. THE LOAD NPC (CS) IN T2 WILL REPLACE THE VALUE SET INTO THE NPC IN T1. THE CONDITIONAL SKIP TESTS THE RESULT OF THE LOGICAL "AND" OPERATION BETWEEN THE KS FIELD OF THE NANO-PROGRAM AND FIST. THE SKIP WILL TAKE EFFECT IF THE RESULT OF THE "AND" IS "NOT" ZERO.

AT T3 THE MPC WILL BE INCREMENTED ONCE, UNCONDITIONALLY. IF T2 HAD NOT BEEN SKIPPED THE MPC WILL BE 2 GREATER THAN AT THE BEGINNING OF THE SKIP-ON-PLUS INSTRUCTION. T3 AND T4 COMPLETE THE ACCESS TO THE APPROPRIATE NANOWORD, AS SELECTED DURING T1 AND T2.

7.5 ARITHMETIC NANOPROGRAMS (SHIFTS)

THIS SECTION WILL ILLUSTRATE SEVERAL EXAMPLES OF ALU, INDEX ALU, SHIFTER, AND F REGISTER INCREMENT / DECREMENT FACILITIES. ALTHOUGH IT IS NOT SHOWN IN ANY ONE EXAMPLE, ALL OF THESE FUNCTIONAL UNITS MAY BE USED SIMULTANEOUSLY AND INDEPENDENTLY. ONLY THE STANDARD ALU AND SHIFTER MAY BE USED FOR COMBINED FUNCTIONS, IN SUPPORT OF DOUBLE PRECISION SHIFTS AND RELATED ALU / SHIFT OPERATIONS.

THE INDEPENDENCE OF THE ALU AND SHIFTER MAY BE SHOWN IN THE SNAP INSTRUCTION. SNAP SIMULTANEOUSLY EXCHANGES THE CONTENTS OF THE LOCAL STORE REGISTER IDENTIFIED BY THE A FIELD WITH THE REGISTER IDENTIFIED BY THE B FIELD. THE ALU IS USED TO TRANSFER R(B) TO R(A) WHILE THE SHIFTER IS USED TO TRANSFER R(A) TO R(B). THE ALU OPERATION COMMAND IS PASS LEFT, WHICH USES ONLY THE LEFT ALU INPUT AND TRANSMITS THE INPUT DATA TO THE ADD BUS WITHOUT MODIFICATION.

```

        SNAP  A,B      [ R(A)=R(B); R(B)=R(A) ]
SNAP:   "SNAP THE CONTENT OF THE A AND B LOCAL STORE REGISTERS"
....   FETCH, KALC = PASS LEFT
X...   A->FSID, B->FSOD, B->FAL
.X...   A->FROD
..X.   GATE SH, GATE ALU
...X

```

NO KSHC FIELD DECLARATION WAS NEEDED AS THE DEFAULT SHIFTER CONDITION (0) IS SINGLE, LEFT, CIRCULAR. A ZERO SHIFT AMOUNT IS ALSO THE KSHA FIELD DEFAULT, AND WILL ALWAYS RESULT IN AN UNMODIFIED SHIFTER OUTPUT. T1 IS USED TO SET UP BOTH INPUT F REGISTERS, IN ORDER TO BEGIN THE PROPAGATION OF THE DATA THROUGH THE FUNCTIONAL UNITS AS EARLY AS POSSIBLE. TWO T-PERIODS ARE REQUIRED BEFORE THE DATA IS STABLE ON THE OUTPUT BUSES. THE OUTPUT F VALUES ARE SET DURING T1 AND T2, THE ONLY REQUIREMENT BEING THAT THEY BE SET PRIOR TO THE T-STEP THAT GATES THEIR OUTPUTS. THE ALU AND SHIFTER OUTPUTS ARE STABLE AT THE END OF T3, WHERE BOTH ARE SIMULTANEOUSLY GATED. IF A LOOK-AHEAD MICROINSTRUCTION FETCH WAS IN USE, IN PLACE OF FETCH, THE ENTIRE INSTRUCTION COULD COMPLETE EXECUTION IN ONLY 3 T-PERIODS.

THE ALU MAY BE USED FOR A FULL ASSORTMENT OF ARITHMETIC AND LOGICAL OPERATIONS. MOST ARITHMETIC OPERATIONS REQUIRE USE OF CARRY-IN AND CARRY-OUT LOGIC, WHILE LOGICAL OPERATIONS NORMALLY DO NOT. IN STANDARD 18 BIT, TWO'S COMPLEMENT ADDITION AND SUBTRACTION THE ALU CARRY-IN MUST BE CLEARED OR SET, RESPECTIVELY. USE OF THE OPPOSITE CARRY-IN WILL RESULT IN A VALUE ONE GREATER THAN THE

CORRECT RESULT FOR ADD OPERATIONS, AND ONE LESS FOR SUBTRACT OPERATIONS. THIS MAY BE USED TO ADVANTAGE IN MULTIPLE PRECISION OPERATIONS WHERE LOW ORDER ADDITIVE CARRY-OUTS, OR SUBTRACTIVE BORROWS, REQUIRE CORRECTIVE ACTIONS IN THE HIGHER ORDER RESULTS.

THE NEXT EXAMPLE IS A SIMPLE SINGLE PRECISION ADD. THE CARRY-IN HOLD MUST BE CLEARED BEFORE THE ALU WILL BEGIN TO COMPUTE A CORRECT TWO'S COMPLEMENT ADDITION.

```

      ADD  A,B      [ R(A)=R(A)+R(B) ]
ADD:    "ADD LOCAL STORE REGISTER A TO B, RESULT IN A"
....   FETCH, KALC = ADD,  ALU STATUS ENABLE
X...   A->FAIL, B->FAIR, A->FAOD, CLEAR CIH
.X...
..X... GATE ALU, ALU TO COH
...X

```

ALL INPUTS ARE DEFINED, INCLUDING THE CARRY-IN-HOLD, IN T1. THE ALU RESULTS ARE AVAILABLE FOR GATING AT THE END OF T3. THE K-VECTOR BIT "ALU STATUS ENABLE" WILL NORMALLY BE USED IN ARITHMETIC AND LOGICAL MICROINSTRUCTIONS, PERMITTING THE ALU STATUS BITS IN FIST TO BE UPDATED UPON GATE ALU. THIS ALLOWS FUTURE INSTRUCTIONS TO TEST THE RESULTS OF THIS OPERATION. THE ALU CARRY-OUT MUST BE MANUALLY TRANSFERRED TO FIST THROUGH THE ALU TO COH COMMAND, WHICH ALSO SAVES THAT CONDITION BIT IN THE CARRY-OUT-HOLD REGISTER. IF THIS INSTRUCTION IS A LOW ORDER COMPONENT OF A MULTIPLE PRECISION ADD THEN CARRY-OUT EQUAL TO 1 WILL INDICATE THAT THE RESULT WAS A 19 BIT VALUE, AND WILL REQUIRE AN ADDITION OF 1 TO THE NEXT HIGHER ORDER ELEMENT. AS STATED ABOVE, THIS EXTRA 1 ADDITION MAY BE ACCOMPLISHED DURING THE NEXT ADD OPERATION BY SETTING THE CARRY-IN-HOLD.

THE NEXT EXAMPLE ILLUSTRATES A DOUBLE PRECISION SUBTRACT INSTRUCTION. DSUB WILL SUBTRACT ONE PAIR OF LOCAL STORE REGISTERS FROM ANOTHER. UNLIKE SOME CONVENTIONAL DOUBLE PRECISION INSTRUCTIONS, WHERE THE REGISTER NUMBERS IN THE INSTRUCTION OPERANDS POINT TO THE HIGH ORDER REGISTERS, OUR INSTRUCTION OPERANDS POINT TO THE LOW ORDER REGISTERS (FOR SIMPLIFICATION OF EXAMPLE). IN THIS EXAMPLE THE PROPAGATION OF THE CARRY-OUT FROM THE LOWER PRECISION RESULT TO THE CARRY-IN OF THE HIGHER PRECISION OPERATION IS ACCOMPLISHED THROUGH AN INTERNAL CARRY COMMAND, "ALU TO BOTH CIH AND COH".

```

        DSUB  A,B      [ R(A-1).R(A)=R(A-1).R(A)-R(B-1).R(B) ]
DSUB:   "SUBTRACT REGISTER PAIR B FROM REGISTER PAIR A"
....   FETCH, ALU STATUS ENABLE, KALC = SUB, KA = RIGHT CTL
X...   A->FAIL,    B->FAIR,    A->FAOD,    SET CIH
.S...   DECF->FAIL, DECF->FAIR, A->FSID,    GATE ALU,    ALU TO BOTH CIH AND COH
..S...   KA->KSHC,    DECF->FAOD
...X                                     GATE ALU,    ALU TO COH

```

T1 SETS ALL ALU LOCAL STORE POINTERS AND THE INITIAL CARRY-IN-HOLD AS REQUIRED FOR A SUBTRACT OPERATION. THE LOW ORDER SUBTRACT PROCEEDS DURING T2. AT THE END OF T2 THE LOW ORDER RESULT IS GATED, THE CARRY-OUT REQUIRED FOR THE HIGHER ORDER SUBTRACT IS PASSED BACK TO THE CARRY-IN-HOLD, AND THE ALU INPUTS ARE DECREMENTED BY ONE TO POINT TO THE HIGH ORDER DATA LOCAL STORE REGISTERS. T3 IS NOW USED TO DECREMENT THE ALU OUTPUT POINTER, AND TO SET THE SHIFTER CONTROL TO RIGHT-CONTROL MODE. THIS SHIFT MODE, ALONG WITH FSID POINTING TO THE LOW ORDER RESULT REGISTER (SET IN T2), WILL ENABLE PROPER DETECTION OF THE 35 BIT "RESULT ZERO" CONDITION. FINALLY T4 GATES THE HIGH ORDER RESULT AND SETS THE LAST CARRY-OUT INDICATION INTO THE CARRY-OUT-HOLD. UPON COMPLETION OF THIS NANOPROGRAM THE CARRY, OVERFLOW, AND SIGN BITS OF FIST ACCURATELY DEPICT THE FINAL VALUE OF THE DOUBLE PRECISION OPERATION. THE FIST "RESULT" BIT SHOWS THE CONDITION OF THE LOW ORDER 35 BITS OF THE AOD AND SOD BUSES, AS DETECTED BY RECYCLING THE LOW ORDER RESULT THROUGH THE SHIFTER DURING THE GATE ALU.

THE INDEX ALU OPERATES WITHOUT THE REQUIREMENT FOR ANY ASSOCIATED F REGISTER TO BE PRESET. ALL INDEX ACTIONS MUST BE STATED, AND COMPLETED, WITHIN ONE STRETCHED T-STEP. THE FOLLOWING EXAMPLE MODIFIES THE LDCSR INSTRUCTION, SEE SECTION 7.3, TO ACCESS A CONTROL STORE REGISTER RELATIVE TO A BASE ADDRESS INSTEAD OF USING AN ABSOLUTE ADDRESS. THE BASE ADDRESS IS MAINTAINED IN AN X REGISTER NAMED X.BASE.

```

        LDCSRB A,B      [ R(A)=CS(B+X(BASE)) ]
LDCSRB: "LOAD REGISTER FROM CONTROL STORE REGISTER PLUS BASE"
....   FETCH, KB = X BASE, KA = 31.
X...   A->FCOD, KA->FCIA, FIPH->A
.S...   INDEX (AOD), INDEX REG (KB) = KA
...X
...S   READ CS (CIA), GATE CS

```

THE FINAL ADDRESS CONSISTS OF THE ORIGINAL B FIELD CONTENT PLUS THE BASE ADDRESS VALUE FROM X.BASE. THIS ADDRESS IS READY AT THE END OF T2, AND MAY BE USED BY THE READ CS NANOPRIMITIVE AT THE BEGINNING OF T4. IN THIS EXAMPLE THE

ADD OPERATION IS COMMANDED DIRECTLY WITHIN THE FSEL2 FIELD OF T-VECTOR T2. OPTIONALLY, THE FSEL2 FIELD MAY BE USED AS A POINTER TO A COMMAND IN ONE OF THE F REGISTERS. THE INDEX REGISTER IS POINTED TO BY THE CONTENT OF THE KB FIELD OF THE K-VECTOR, AND THE LOCAL STORE REGISTER IS POINTED TO BY THE KA FIELD. THESE REGISTER SELECTIONS MAY ONLY BE MADE THROUGH INDIRECT POINTERS, UNLIKE THE COMMAND SELECTION FIELD.

THE SHIFTER MAY BE USED FOR EITHER SINGLE LENGTH (18 BIT) OR DOUBLE LENGTH (36 BIT) SHIFT OPERATIONS. WHEN USED FOR SINGLE LENGTH SHIFTS ONLY THE SID AND SOD BUSES ARE USED, WITH THE KSHC AND KSHA FIELDS SPECIFYING THE TYPE OF SHIFT AND SHIFT AMOUNT, RESPECTIVELY. DOUBLE LENGTH SHIFTS REQUIRE USE OF THE SID BUS AND ONE OR BOTH ALU BUSES (AIL AND AIR) FOR INPUT, AND THE SOD AND AOD BUSES FOR OUTPUT. IF THE ALU IS NOT USED FOR A COMBINED ALU FUNCTION AND SHIFT OPERATION IT MAY ESSENTIALLY BE BYPASSED THROUGH USE OF THE PASS LEFT ALU FUNCTION. THIS ALLOWS THE DOUBLE LENGTH SHIFT TO COMPLETE IN 2 T-PERIODS. USE OF MOST OTHER ALU COMMANDS REQUIRES AN EXTRA T-PERIOD FOR COMPLETION OF THE ALU PHASE OF THE OPERATION, BEFORE ITS OUTPUT IS READY FOR THE HIGH ORDER ACTION OF THE SHIFTER. THE FOLLOWING EXAMPLE PERFORMS A DOUBLE LENGTH SHIFT, WHERE THE ALU IS USED STRICTLY TO PASS THE HIGH ORDER 18 BITS TO THE SHIFTER EXTENSION.

```

SRDAI  A,B      [ R(A),R(A+1)=R(A),R(A+1)->B ]
SRDAI:  "SHIFT RIGHT DOUBLE ARITHMETIC IMMEDIATE"
....   FETCH, KSHC=RIGHT+DOUBLE+ARITHMETIC+RIGHT CTL, SH STATUS ENABLE
        KALC=PASS LEFT,                               ALU STATUS ENABLE
S...   A->FAIL,   A->FSID,   B->KSHA, CLEAR CIH
.S...   A->FAOD, INCF->FSID,   A->FSOD
..S...                               INCF->FSOD
...X   GATE ALU,   GATE SH, ALU TO COH

```

IN A SIMILAR FASHION TO THE ADD EXAMPLE ABOVE, THE K-VECTOR SH STATUS ENABLE BIT ALLOWS FIST TO BE UPDATED AT GATE SH TIME WITH THE SHIFTER HIGH BIT AND LOW BIT STATUS. THESE BITS THOUGH ARE STRICTLY THE OUTPUTS OF BITS 0 AND 17 ON THE SOD BUS. THE ALU STATUS ENABLE ALLOWS THE SETTING OF THE SIGN, CARRY, OVERFLOW, AND RESULT BITS OF FIST, UPON GATE ALU EXECUTION. IN THE CASE OF A RIGHT SHIFT OPERATION OVERFLOW WILL ALWAYS BE ZERO. CARRY SHOULD ALSO BE ZERO, AND THE CLEAR CIH OPERATION IN T1 WILL GUARANTEE THIS. SIGN WILL BE SET TO THE CORRECT 36 BIT RESULTANT SIGN. THE RESULT BIT OF FIST WILL BE SET ZERO IF ALL OF THE LOW ORDER 35 BITS (0 THROUGH 34) ON THE AOD AND SOD BUSES ARE ZERO, AS DETECTED BY "RIGHT CTL" MODE.

THIS INSTRUCTION IS DEFINED TO SHIFT THE REGISTERS A AND A+1. THEREFORE THE LOW ORDER SHIFT OPERATIONS WILL REQUIRE THAT FSID AND FSOD BE SET ONE GREATER THAN THE VALUE IN THE A FIELD. THIS IS ACCOMPLISHED BY TRANSFERRING THE VALUE OF A TO ALL REQUIRED F REGISTERS, AND THEN INCREMENTING THAT VALUE DIRECTLY IN FSID AND FSOD WITH THE 6 BIT INCREMENT FUNCTIONS IN T2 AND T3 RESPECTIVELY. BOTH INPUTS ARE STABLE IMMEDIATELY AFTER T2, WHICH WOULD ALLOW THE RESULTS TO BE GATED IN A STRETCHED T3 OR AT ANY TIME AFTER T3. FASTER TECHNIQUES FOR SETTING UP THE FINAL F REGISTER VALUES MAY BE WORKED OUT, WITH THE INDEX ALU FOR EXAMPLE, REDUCING THE OVERALL TIME REQUIRE FOR THE SRDAI INSTRUCTION.

AS DISCUSSED ABOVE, WHEN THE ALU AND SHIFTER ARE USED FOR A COMBINED FUNCTION THE TIME REQUIRED FOR THE PASSAGE OF DATA THROUGH BOTH THE ALU AND SHIFTER EXTENSION IS 3 T-PERIODS. AN EXAMPLE OF A COMBINED OPERATION WHERE THIS CAPABILITY MAY BE USED IS IN THE EXTRACTION OF SUBFIELDS OF DATA WORDS. TO EXTRACT A SUBFIELD AND RIGHT JUSTIFY THE EXTRACTED RESULT IN A REGISTER WE NEED SEVERAL PIECES OF INFORMATION. FIRST THE ACTUAL DATA SOURCE AND THE DESTINATION REGISTERS MUST BE IDENTIFIED. SECOND THE ACTUAL SUBFIELD WIDTH AND POSITION MUST BE DESCRIBED IN SOME FORM. FOR OUR NEXT EXAMPLE WE WILL SIMPLIFY THESE ITEMS AS FOLLOWS. THE SOURCE AND DESTINATION DATA REGISTERS WILL BE THE SAME, AS IDENTIFIED BY THE A FIELD OF OUR MICROINSTRUCTION. THE ACTUAL ORIGINAL FIELD WIDTH AND SITUATION WILL BE FOUND AS A PREDEFINED MASK IN THE LOCAL STORE REGISTER SYMBOLICALLY NAMED LS.MASK. FINALLY, THE RIGHT JUSTIFICATION WILL BE DESCRIBED BY A SHIFT AMOUNT SPECIFIED AS THE MICRO-INSTRUCTION B FIELD.

```

      MASK.SHIFT  A,B      [ R(A)=(R(A)&R(LS.MASK))>>B ]
MASK.SHIFT: "MASKED-SHIFT REGISTER"
....  FETCH, KSHC = RIGHT+DOUBLE+LOGICAL, KALC = AND, KA = LS.MASK
S...  KA->FAIR, A->FAIR, B->KSHA, CLEAR COH
.X...  A->FAOD
...X.
...X  GATE ALU

```

ALTHOUGH THE SHIFTER IS ACTUALLY INVOLVED IN OUR INSTRUCTION OPERATION ONLY THE HIGH ORDER INPUT AND OUTPUT COMPONENT AFFECT OUR RESULT. THEREFORE FSID AND FSOD ARE NOT REFERENCED, NOR IS A GATE SH NEEDED. THE ALU INPUTS ARE SET UP IN T1. THE RIGHT INPUT POINTS TO THE FIXED REGISTER LS.MASK AND THE LEFT INPUT POINTS TO THE A FIELD SELECTED INPUT. THE SHIFT AMOUNT IS ALSO SET IN T1, PERMITTING THE ENTIRE OPERATION TO BEGIN IMMEDIATELY FOLLOWING T1. THE CLEAR COH PRIMITIVE IS NECESSARY TO AVOID THE PROPAGATION OF AN EXTRA HIGH ORDER BIT INTO THE RESULT DURING THE LOGICAL SHIFT OPERATION. THE ALU OUTPUT SETTING MAY

BE MADE ANY TIME PRECEDING THE GATE ALU. THE ALU FUNCTIONS AS FOLLOWS. FIRST THE LOGICAL PRODUCT OF THE CONTENT OF REGISTER A AND THE MASK REGISTER IS COMPUTED. THEN THAT RESULT IS SHIFTED RIGHT, WITH LEFT ZERO FILL, AND PLACED ONTO THE ADD BUS. THIS NORMALLY TWO STEP OPERATION IS NOW COMPLETE IN ONE EXTENDED ALU-SHIFTER OPERATION.

7.6 MULTINORD NANOPROGRAMS

WHEN THE ALGORITHM FOR A NANOPROGRAM REQUIRES MORE THAN 4 T-STEPS ADDITIONAL NANOWORDS MUST BE USED. THERE ARE SEVERAL ALTERNATIVE METHODS THAT MAY BE USED IN TRANSFERRING CONTROL BETWEEN NANOWORDS. SELECTION OF THE APPROPRIATE METHOD REQUIRES KNOWLEDGE OF THE PROGRAMMING CONVENTIONS IN USE WITHIN THE SET OF RESIDENT NANOPROGRAMS MAKING UP THE ACTIVE NANOSTORE ENVIRONMENT. THE MOST SIMPLE NANOSTORE ENVIRONMENT WOULD BE ONE WHERE THERE ARE NO NANOWORDS SHARED BY DIFFERENT NANOPROGRAMS. IN MOST NANOPROGRAMMING THE PROGRAMMER TRIES TO IDENTIFY ALL COMMON SETS OF PROCEDURES. IF THERE EXISTS FREQUENTLY REPEATED EXIT CODE, FROM SEVERAL NANOPROGRAMS, IT MAY BE POSSIBLE FOR ALL TO USE THE SAME LAST NANOWORDS (REFERRED TO IN THE FUTURE AS COMMON TAIL NANOWORDS). IN A SIMILAR MANNER, IF THERE EXISTS FREQUENTLY REPEATED SETS OF NANOCODE WITHIN DIFFERENT NANOPROGRAMS THEN IT MAY BE POSSIBLE TO SET UP A COMMON BODY NANO-SUBROUTINE.

INITIAL ENTRY INTO THE FIRST NANOWORD OF A MICROINSTRUCTION INITIATED NANOPROGRAM USES "LOAD NPC(CS)" TO SET THE ENTRY ADDRESS INTO THE NANOPROGRAM COUNTER. THE NPC MAY BE MODIFIED FROM OTHER SOURCES. THE KN FIELD OF THE CURRENT NANOWORD MAY BE TRANSFERRED INTO NPC, USING LOAD NPC(KN). NPC MAY BE INCREMENTED BY ONE, FOR EACH USE OF LOAD NPC(SEQ). CONTROL MAY ALSO BE TRANSFERRED WITHOUT MODIFYING THE CURRENT NPC, USING THE NANOBANCH SELECTION MECHANISM. IN THIS CASE THE KN FIELD IS USED DIRECTLY TO READ THE NEXT NANOSTORE LOCATION.

AN EXAMPLE OF THE ENTRY / EXIT PROCEDURE FOR CALLING ON A NANO-SUBROUTINE MAY BE TO REQUIRE THE CALLING NANOPROGRAM TO SET THE RETURN NANOSTORE ADDRESS INTO THE NPC PRIOR TO ENTERING THE SUBROUTINE. THIS MAY BE ACCOMPLISHED BY EITHER INCREMENTING THE NPC, WITH LOAD NPC(SEQ), OR BY ACTUALLY SETTING UP THE ENTRY ADDRESS OF THE NEXT NANOPROGRAM USING THE NEXT MICROINSTRUCTION OPERATION CODE, WITH LOAD NPC(CS). TRANSFER TO THE SUBROUTINE WOULD THEN BE VIA NANOBANCH ADDRESS SELECTION, WITH THE BRANCH(ADDRESS) NANOPRIMITIVE. THE SUBROUTINE ITSELF MAY TRANSFER TO MULTIPLE NANOWORDS USING ONLY THE NANOBANCH MECHANISM. THIS SUBROUTINE WOULD EXIT BY SIMPLY READING NANOSTORE WITHOUT ANY ALTERATION OF NPC. THE LOCATION READ WILL BE THAT ORIGINALLY DESIRED FOR THE RETURN, BY THE CALLING NANOPROGRAM. FOR EXAMPLE:

CALLER	[]	SUBROUTINE
	[]	
BRANCH(SUBROUTINE)	[]	
T(N) LOAD NPC(SEQ)	[]	T(M) READ NS, GATE NS
T(N+1) READ NS, GATE NS	[]	"RETURNS TO CALLER + 1"

THE FOLLOWING EXAMPLE DEFINES THE BALTNW (BRANCH AND LINK ON TRUE TO NEXT WORD) INSTRUCTION. THIS INSTRUCTION REQUIRES TWO NANOWORDS. THE TECHNIQUE USED FOR NANOSTORE ADDRESS MODIFICATION IS INCREMENT NPC. THE BRANCH AND LINK OPERATION TAKES PLACE ONLY WHEN AT LEAST ONE OF THE BITS SELECTED IN THE B FIELD IS ALSO SET IN FIST. IF NONE OF THE SELECTED BITS ARE SET CONTROL PROCEEDS TO THE NEXT SEQUENTIAL INSTRUCTION.

```

        BALTNW  A,V,B
    I IF ( B .AND. FIST ) THEN R(A)=MPC+2, MPC=V, ELSE MPC=MPC+2 J
BALTNW: "BRANCH AND LINK TO CS(V) ON TRUE"
....  LEGAL MICRO OP ENTRY, KB = LS.MPC
S...  B->KS,          KB->FSID,          MPC PLUS 2,          LOAD NPC (SEQ)
.S...  KB->FCOD,      A->FSOD,          READ CS (MPC+1), READ NS,  GATE NS (NOT S)
.S...  READ CS (NPC), LOAD NPC (CS)
.S...  READ NS,      GATE NS,          LOAD R31

:
....  ALLOW INTS
X...  GATE SH,      GATE CS
.S...  READ CS (COD), LOAD NPC (CS)
.S...  READ NS,      GATE NS,          LOAD R31

```

IN THE FIRST WORD, T1 AND T2 SET UP THE "S" TEST CONDITION FROM THE B FIELD AND EXECUTE THE TEST. IF THE RESULT INDICATES "FALSE" THEN CONTROL STAYS WITHIN THIS WORD, FETCHING THE NEXT SEQUENTIAL MICROINSTRUCTION. OTHERWISE, CONTROL TRANSFERS TO THE NEXT SEQUENTIAL NANOWORD WHERE THE RETURN MPC ADDRESS IS SAVED IN R(A) AND MPC IS SET TO THE VALUE OF THE SECOND NORD OF THE BALTNW INSTRUCTION ITSELF. THE SHIFTER IS USED TO PASS THE OLD MPC TO R(A). CONTROL STORE OUTPUT DATA BUS IS USED TO SUPPLY THE ADDRESS FOR THE READ CS IN T2 OF WORD 2, AND FOR THE NEW MPC VALUE VIA GATE CS IN T1 OF WORD 2.

THE FIRST READ NS IN WORD 1 READS WORD 2. THIS ADDRESS IS SELECTED IN T1 VIA THE LOAD NPC (SEQ). IF NOT USED, NPC IS RELOADED FROM THE MICROINSTRUCTION SELECTED NANOSTORE ADDRESS IN T3.

NOTE THAT ALTHOUGH IT APPEARS THAT THE ORIGINAL MPC VALUE IS INCREMENTED BY 2, IN T1 OF WORD 1, THE READ CS (MPC+1) IN T2 STILL USES THE ORIGINAL MPC VALUE. THIS OCCURS BECAUSE MPC PLUS 2 IS A TRAILING EDGE OPERATION AND READ CS IS A LEADING EDGE OPERATION. DATA PROPAGATION DELAY TIMES GUARANTEE THAT THE MPC VALUE CANNOT CHANGE BEFORE THE CONTROL STORE ADDRESS IS COMPLETELY DECODED.

```

                                ADD/PASS TEST BIT
                                ^\
PRODUCT OUTPUT (AOD)          PRODUCT///MULTIPLIER ""
                                OUTPUT (SOD)  ""
=====
!CARRY! !                      ! !
! OUT !-->! SHIFTER LEFT HALF !-->! SHIFTER RIGHT HALF ""!
!HOLD.! !                      ! !
=====
      // // // // // /---/
      // // // // // /---/ (SID) OFFSET MULTIPLIER INPUT
=====
! CARRY! !
!<-----! A L U (ADD OR PASS)!
  OUT !
=====
PRODUCT INPUT (AIL) : (AIR) MULTIPLICAND INPUT

```

THE ABOVE ILLUSTRATION SHOWS THE ACTUAL ALU-SHIFTER CONNECTIONS DURING THE SECOND WORD EXECUTION, UMULT2, ONLY. FOLLOWING EACH ADDITION OPERATION THE ALU CARRY-OUT IS PLACED INTO THE CARRY-OUT-HOLD REGISTER. THIS CARRY VALUE IS THEN TRANSFERED INTO THE AOD HIGH BIT DURING THE RIGHT SHIFT BY 1 OPERATION. THE DECISION WHETHER TO ADD THE MULTIPLICAND INTO THE CURRENT RESULT OR ONLY TO PASS THE CURRENT RESULT THROUGH THE ALU TO THE SHIFTER IS MADE BY TESTING THE CARRY-OUT-HOLD REGISTER VALUE AFTER THE SHIFT OPERATION. THIS VALUE WILL BE SET FROM THE VALUE OF THE SHIFTER LOW BIT (SOD BIT 0) AT THE END OF EACH CYCLE.

```

        UMULT  A,B [ R(A),R(B)=R(A)*R(B) ]
UMULT:  "MULTIPLY A TIMES B"
....   LEGAL MICRO OP ENTRY, BRANCH (UMULT2), KA = ZERO, KB = LS.WORK,
        KT = 18., KX = PASS LEFT, KALC = PASS LEFT,
        KSHC = RIGHT+SINGLE+LOGICAL+RIGHT CTL
X...   A->FAIL, KB->FAOD, B->FSID, CLEAR CIH
.S...   KB->FAIR, A->FAOD, B->FSOD, GATE ALU, READ CS (MPC+1), MPC PLUS 1,
                                         LOAD NPC (CS)
.S...   KA->KALC, KX->F.PASS, SH TO COH, READ NS
.S...   KT->F.COUNT, GATE ALU, GATE NS

: UMULT2 = N. "CONTINUATION OF UNSIGNED MULTIPLY"
....   BRANCH (N.+1), KALC = PASS LEFT, KSHC = RIGHT+DOUBLE+LOGICAL, KSHA = 1,
        KT = CARRY, KX = F ZERO, KB = ADD
.S...   KB->KALC, SKIP (NOT T)
.X...   F.PASS->KALC
.S...   ALU TO COH, READ NS, GATE NS (X), DECF->F.COUNT
.X...   SH TO COH, GATE ALU, GATE SH

: "COMPLETION OF UNSIGNED MULTIPLY"
....   ALLOW INTS, KSHC = LEFT+DOUBLE+LOGICAL+RIGHT CTL, KALC = PASS LEFT
        ALU STATUS ENABLE
X...
.S...   ALU TO COH, GATE ALU, GATE SH, READ NS, GATE NS, LOAD R31

```

THE FOLLOWING DISCUSSION WILL COVER THE ACTION OF EACH NANOWORD IN THE ABOVE EXAMPLE. IN THE INITIALIZATION WORD, T1 PREPARES THE ALU TO PASS THE MULTIPLICAND TO LOCAL STORE REGISTER LS.WORK FOR USE AS THE MULTIPLICAND SOURCE DURING ADDITIONS. CARRY-IN-HOLD IS ALSO CLEARED FOR THE ADD OPERATIONS THAT WILL FOLLOW. T2 EXECUTES THE GATE ALU, SAVING THE MULTIPLICAND. THE ALU AND SHIFTER BUS CONNECTIONS ARE COMPLETED. ALL MICROINSTRUCTION FETCH ACTIONS ARE COMPLETED: READ CS (MPC+1), LOAD NPC (CS), AND MPC PLUS 1. THE NEXT NANOPROGRAM ADDRESS IS NOW SELECTED, BUT WILL NOT BE REFERENCED UNTIL THE READ NS IN THE COMPLETION WORD. T3 CHANGES THE ALU FUNCTION TO ZERO, WHICH WILL CLEAR THE INITIAL PRODUCT VALUE IN R(A). AN F REGISTER NAMED F.PASS IS INITIALIZED WITH THE VALUE OF THE ALU PASS FUNCTION. COH IS SET TO THE VALUE OF THE RIGHT MOST BIT OF THE MULTIPLIER, TAKEN FROM SID BIT 0 USING THE "RIGHT CTL" FUNCTION OF THE SHIFTER. NANOSTORE LOCATION UMULT2 IS READ. FINALLY, T4 SETS AN F REGISTER NAMED F.COUNT TO THE VALUE 18 (DECIMAL), WHICH WILL BE USED BY UMULT2 AS A COUNTER DURING ITS 18 LOOPS. REGISTER R(A) IS ZEROED VIA THE ALU, AND CONTROL IS TRANSFERED TO T1 OF UMULT2.

UMULT2 PERFORMS THE ACTUAL MULTIPLICATION. THIS REQUIRES IT TO REPEAT ITS FULL 4 T-STEPS 18 TIMES. T1 IS USED TO SET THE ALU CONTROL TO ADD MODE, AND TO DETERMINE WHETHER TO CHANGE THE FUNCTION TO PASS LEFT. IF THE CURRENT RIGHT-MOST BIT OF THE MULTIPLIER IS A 1 THEN T2 WILL BE SKIPPED, LEAVING THE ALU SET FOR AN ADD OPERATION. IF A 0 THEN T2 IS EXECUTED AND THE ALU WILL NOT ALTER THE RESULT DURING THIS LOOP CYCLE. T3 TRANSFERS THE ALU CARRY-OUT CONDITION TO THE CARRY-OUT-HOLD, AS REQUIRED BEFORE THE ACTUAL GATE ALU IS PERFORMED, TO PRESET CARRY-OUT-HOLD WITH THE CORRECT VALUE TO BE PROPAGATED INTO THE SIGN BIT POSITION OF ADD. T3 ALSO READS THE NEXT NANOSTORE LOCATION, AND MAKES THE DECISION TO TERMINATE THE MULTIPLICATION WHEN F.COUNT IS DECREMENTED TO ZERO. T4 NOW COMPLETES THE CYCLE BY GATING THE NEW PARTIAL PRODUCT INTO R(A) AND R(B) ALONG WITH THE SHIFTING OF THE MULTIPLIER RIGHT ONE BIT POSITION. THE SH TO COH OPERATION SETS COH TO THE VALUE OF THE NEW RIGHT-MOST BIT ON SOD, TO BE USED FOR THE ADD/PASS DECISION IN T1 FOR THE NEXT LOOP CYCLE.

THE COMPLETION WORD IS REQUIRED ONLY IF IT IS DESIRED TO SET FIST TO ACCURATELY PORTRAY THE SIGN AND RESULT OF THE FINAL 36 BIT PRODUCT. T1 IS EMPTY, ALLOWING FOR INTERRUPT ADDRESS SELECTION AND FOR ALU-SHIFTER PROPAGATION. T2 GATES THE ALU AND SHIFTER BACK INTO THEIR CURRENT REGISTERS, UNMODIFIED, ONLY TO CAUSE THE CORRECT SETTING OF SIGN, RESULT, OVERFLOW AND CARRY IN FIST. CARRY AND OVERFLOW ARE MEANINGLESS IN AN UNSIGNED MULTIPLY OPERATION.

SINCE UMULT2 IS ACTUALLY A COMPLETE MULTIPLY ROUTINE ANY NANOPROGRAM REQUIRING A MULTIPLICATION OPERATION AS ITS LAST PROCEDURE MAY USE IT AS A COMMON TAIL. WITH MINOR ALTERATIONS, UMULT2 MAY ALSO BE USED AS THE FINAL PHASE OF A SIGNED MULTIPLY ROUTINE.

8 QM-1 I/O SYSTEMS

8.1 GENERAL

THE QM-1 HAS 8 INDEPENDENT I/O PORTS (SEE SECTION 4.6, EXTERNAL INTERFACE). ANY OR ALL OF THESE PORTS MAY BE USED FOR CONCURRENT DATA TRANSFERS AND DEVICE CONTROL OPERATIONS. A QM-1 PORT MAY BE INTERFACED DIRECTLY TO A USER'S OWN EQUIPMENT OR TO NANODATA STANDARD PERIPHERAL DEVICES THROUGH THE QM-1 CHANNEL CONTROLLER. REFER TO APPENDIX-A FOR A DESCRIPTION OF ACTUAL QM-1 PORT INTERFACING.

FIGURE 8.1A IS A BLOCK DIAGRAM OF THE STANDARD QM-1 I/O SYSTEM. NO MORE THAN ONE CHANNEL CONTROLLER MAY BE ATTACHED TO EACH CPU PORT. UP TO 64 STANDARD DEVICE CONTROLLERS MAY BE CONNECTED TO EACH CHANNEL CONTROLLER. MORE THAN 64 ACTUAL I/O DEVICES CAN BE PLACED ON A SINGLE CHANNEL DUE TO MANY FORMS OF DEVICE CONTROLLERS SUPPORTING MORE THAN ONE DEVICE (IE. TAPE DRIVES, TELECOMMUNICATIONS DEVICES). FOR CONNECTION OF USER OWNED EQUIPMENT TO A QM-1 CHANNEL CONTROLLER, REFER TO THE DOCUMENT "NANODATA STANDARD CHANNEL CONTROLLER".

DATA TRANSFERS ARE MAINTAINED ON A WORD TO WORD BASIS. ALL DEVICES ON THE SAME CHANNEL MAY BE TRANSFERRING DATA SIMULTANEOUSLY, AS LONG AS THEIR COMBINED DATA RATES DO NOT EXCEED THE DESTINATION MEMORY ACCESS SPEED. UP TO 18 BITS AT A TIME ARE TRANSFERRED BETWEEN THE PORT AND DEVICE, OVER EACH ACTIVE CHANNEL. WHEN DATA IS BEING PASSED DIRECTLY TO A CPU PORT THE CPU IS INTERRUPTED PERIODICALLY, IN ORDER TO ROUTE EACH DATUM BETWEEN THE I/O PORT AND APPROPRIATE QM-1 MEMORY. THE STANDARD QM-1 DEVICE CONTROLLERS MAINTAIN DATA ROUTING INFORMATION FOR THE DURATION OF THE FULL DATA BLOCK TRANSFER. THIS CONSISTS OF AUTOMATIC STORAGE AND UPDATING OF THE MEMORY ADDRESS POINTER AND BLOCK LENGTH WORD COUNTER. IN ADDITION, THE CPU IS NOTIFIED AT THE END OF OPERATION OR OF OTHER PROGRAMMABLY SELECTED CONDITIONS, SUCH AS DEVICE ERRORS, DEVICE READY STATE CHANGE, WORD COUNT REACHING ZERO, ETC.

QM-1 SYSTEMS EQUIPPED WITH THE OPTIONAL MULTIPORT MAIN STORE INTERFACE MAY TAKE ADVANTAGE OF THE DIRECT MEMORY ACCESS (DMA) PATH. WITH THE ADDITION OF THE APPROPRIATE DMA CHANNEL CONTROLLERS, I/O DATA TRANSFERS MAY PROCEED WITHOUT CPU INTERVENTION. IN DMA OPERATING MODE THE CPU IS USUALLY INVOLVED IN THE DATA TRANSFER OPERATION ONLY AT INITIATION, AND FOLLOWING TERMINATION. THIS PERMITS DEVICES WITH HIGH SPEED DATA TRANSMISSION RATES TO BE ACTIVE WITHOUT ADDING ANY DIRECT CPU OVERHEAD. LOW SPEED DEVICES MAY ALSO UTILIZE THE DMA PATH, AT THE INSTALLATION'S OPTION.

THE STANDARD QM-1 MULTIPORT MEMORY CONSISTS OF UP TO 8 EXTERNAL PORTS. AT LEAST ONE OF THESE PORTS MUST BE CONNECTED TO THE NORMAL CPU MAIN STORAGE ACCESS BUS. AN ADDITIONAL PORT IS REQUIRED FOR EACH CPU WITHIN A MULTI-PROCESSOR ENVIRONMENT. ALL REMAINING MAIN STORE PORTS MAY BE CONNECTED TO INDIVIDUAL DMA CONTROLLERS, PERMITTING UP TO SEVEN DMA CONNECTIONS. WHEN MAIN STORE IS CONFIGURED AS A FOUR WAY INTERLEAVED MEMORY (750 NANOSECONDS FULL CYCLE, 18 BIT ACCESS) AN AGGREGATE DATA RATE OF OVER 10 MILLION BYTES PER SECOND MAY BE REALIZED.

THE FOLLOWING SECTIONS DESCRIBE THE INDIVIDUAL COMPONENTS OF THE QM-1 I/O SYSTEM IN FULL DETAIL. FOR A THOROUGH UNDERSTANDING OF QM-1 I/O IT IS RECOMMENDED THAT THE READER ALSO BE FAMILIAR WITH THE QM-1 EXTERNAL INTERRUPT MECHANISM (SECTIONS 4.5.2 AND 4.5.4.2), EXTERNAL STORE (SECTION 4.2.5), AND THE EXTERNAL INTERFACE ORGANIZATION (SECTIONS 4.6 AND 5.4.4).

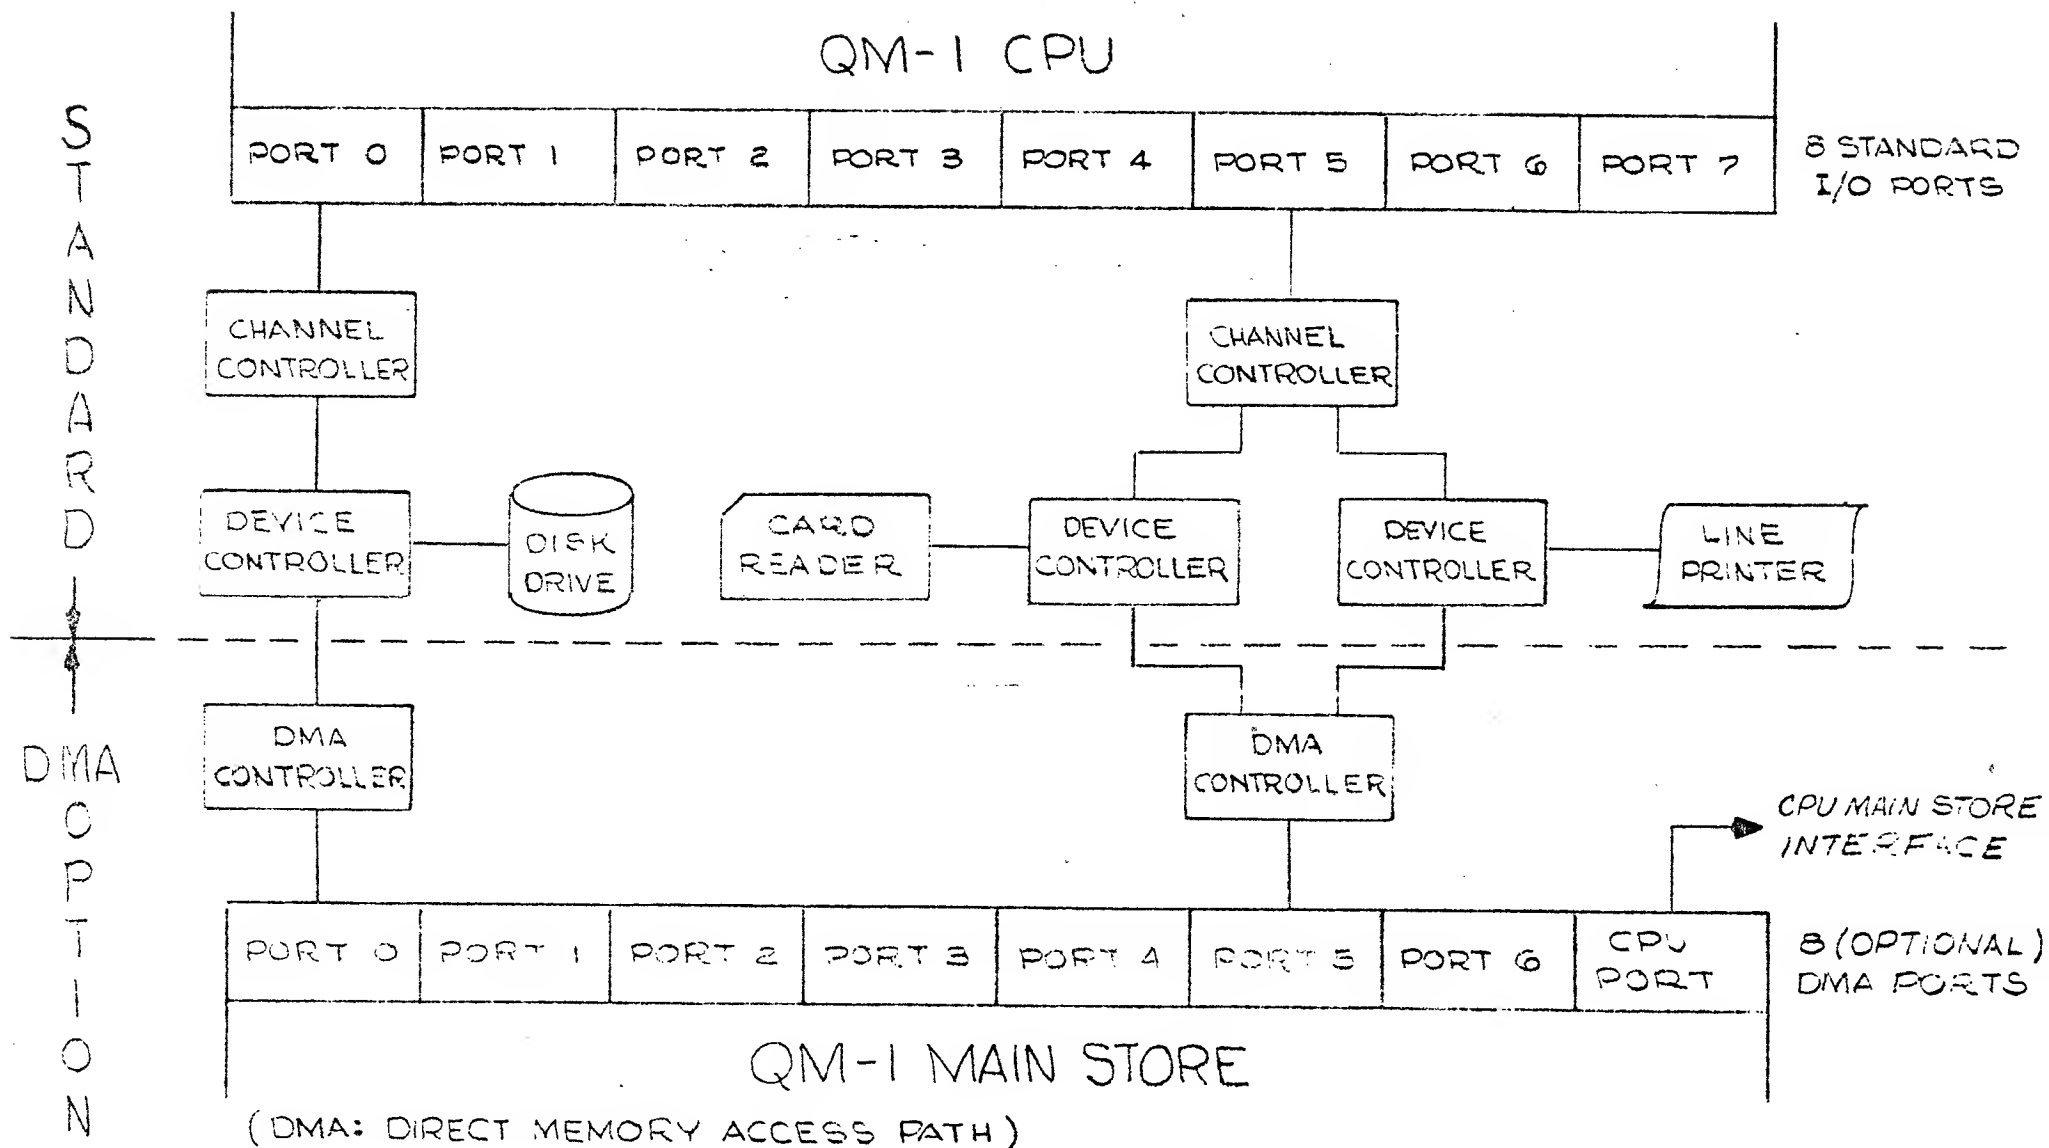


FIG 8.1A - QM-1 STANDARD I/O SYSTEM

8.2 QM-1 I/O OPERATION

8.2.1 GENERAL

THE QM-1'S I/O SYSTEM, BEING CONSISTENT WITH THE DESIGN OF QM-1, CONSISTS OF INDIVIDUAL FUNCTIONS WHICH MAY BE UTILIZED IN NUMEROUS WAYS TO MOVE INFORMATION INTO AND OUT OF THE QM-1 CPU. THE FOLLOWING SECTION IS A DESCRIPTION OF THE RELATIONSHIPS BETWEEN THE VARIOUS I/O SUPPORT FUNCTIONS AND THEIR OVERALL OPERATION RELATIVE TO THE REST OF THE CPU. SECTION 8.2.2 IS CONCERNED WITH THE DATA ROUTING SUPPORT FACILITIES OF QM-1 I/O, AVAILABLE AT THE PORTS, WHILE SECTION 8.2.3 DISCUSSES THE INTERRUPT RELATED OPERATIONS.

8.2.2 QM-1 PORT OPERATION

SECTION 4.6 DESCRIBES A QM-1 PORT AND THE FACILITIES AND CONTROLS AVAILABLE FOR I/O SUPPORT. THIS SECTION IS THEREFORE DEVOTED TO THE RELATIONSHIP OF THESE CONTROLS TO EACH OTHER AND TO OVERALL CPU TIMING.

THE FIRST PIECE OF INFORMATION NECESSARY FOR ANY PORT OPERATION IS "KA". THE LOW THREE BITS OF "KA" SELECT THE PORT AT WHICH AN OPERATION WILL BE PERFORMED. "KA" MUST BE STABLE ON THE LEADING EDGE OF THE T-PERIOD IN WHICH AN I/O OPERATION IS TO TAKE PLACE. THE FOLLOWING PROGRAM SEGMENTS ALL TRY TO SEND AN XIO SIGNAL TO PORT THREE (3). THE FIRST TWO ARE VALID WHILE THE THIRD FAILS.

1. :
 KA = 3 "SELECTS PORT 3"
 X... XIO
2. : "ASSUME FLIV CONTAINS 3"
 X... FLIV->KA
 .X... "WAIT FOR PORT SELECTION"
 ..X... XIO
3. : "ASSUME FLIV CONTAINS 3"
 X... FLIV->KA
 .X... XIO "FAILS BECAUSE KA WAS NOT DEFINED IN TIME"

THESE THREE EXAMPLES HOLD FOR RIO AS WELL.

THE RELATIONSHIP OF LOADING A PORT REGISTER TO AN OPERATION ON THE PORT IS VERY IMPORTANT. THE FOLLOWING PROGRAM SEGMENT GATES MAIN STORE INTO PORT REGISTER 3 (EXTERNAL STORE 3 VIA FMODE = 32 + 3, SEE SECTION 4.2.6) AND THEN SENDS XIO TO PORT 3. THE DATA WILL BE VALID AT THE TIME OF XIO.

```
:  
.... KA = 3 "SELECTS PORT 3", KB = 43 "OCTAL, 35. DECIMAL"  
X... KB->FMODE "POINT MOD TO ES(3)"  
.X.. GATE MS "GATES MOD TO ES(3)"  
..X. XIO "SENDS DATA TO DEVICE"
```

IF XIO HAD BEEN IN T2 INSTEAD OF T3 THE OLD VALUE IN ES(3) WOULD HAVE BEEN AT THE PORT WHEN XIO OCCURRED.

THE SAME TIMING HOLDS WHEN THE EXTERNAL STORE REGISTER IS BEING LOADED FROM LOCAL STORE VIA THE LOAD ES NANOPRIMITIVE.

DATA BEING READ INTO THE CPU HAS MUCH THE SAME TIMING AS THE LOAD ES FUNCTION. THE RIO PRIMITIVE CAUSES THE DATA AT THE "KA" SELECTED PORT TO BE LOADED INTO THE ASSOCIATED PORT REGISTER. THIS FUNCTION IS LEADING EDGE WITH "TRAILING EDGE RESULTS". THIS MEANS THAT THE TRANSFER IS STARTED ON THE LEADING EDGE AND COMPLETED ON THE TRAILING EDGE OF THE T-PERIOD. THIS IS IMPORTANT AS SHOWN IN THE FOLLOWING PROGRAM SEGMENT.

```
: "ASSUME FE0A = 3 AND FE0D = 1"  
.... KA = 3  
S... RIO "READ THE DATA FROM PORT 3 INTO PORT REGISTER 3, ES(3)"  
S... GATE ES "GATE THE SAME DATA INTO LOCAL STORE REGISTER 1"
```

THE GATE ES IS VALID, SINCE THE RIO WAS COMPLETED BY THE TRAILING EDGE OF THE FIRST T-PERIOD OF THE STRETCHED T-STEP. THE GATE ES THEN TOOK PLACE AT THE TRAILING EDGE OF THE SECOND T-PERIOD OF THE SAME T-STEP.

THE 6 BIT TRANSFER PRIMITIVE "IO ID -> F REG" IS A STANDARD TRAILING EDGE FUNCTION WHICH GATES THE CONTENTS OF THE I/O ID LINES OF THE "KA" SELECTED PORT INTO THE SPECIFIED F-REGISTER. TIMING IS THE SAME AS XIO AND RIO WITH RESPECT TO "KA" PORT SELECTION, HOWEVER IT IS A TRUE TRAILING EDGE FUNCTION WHILE XIO AND RIO ARE LEADING EDGE WITH TRAILING EDGE TIMING.

THE TWO SIX BIT BUSES FROM THE QM-1 TO THE PORTS ARE THE "G-BUS" AND THE "PHANTOM-BUS" (SEE SECTION 4.6). ALL FOLLOWING TEXT WILL REFER TO THESE BUSES AS THE "DEVICE SELECTION" AND "CPU COMMAND" BUSES RESPECTIVELY. THESE TERMS REPRESENT SUGGESTED USES OF THESE TWO SIX BIT DATA PATHS, AND IN NO WAY LIMITS THEIR USE FOR ANY OTHER PURPOSE. NANODATA STANDARD QM-1 I/O EQUIPMENT USES THE SIX BITS OF THE "G-BUS" TO SELECT ONE OF SIXTY FOUR (64) DEVICE CONTROLLERS ON A CHANNEL WHILE THE "PHANTOM-BUS" IS USED TO SPECIFY ONE OF SIXTY FOUR (64) COMMANDS TO BE PERFORMED WHEN AN XIO SIGNAL IS SENT TO A CHANNEL.

THE TIMING OF THE COMMAND AND SELECTION LINES IS SUCH THAT THEY ARE STABLE AT XIO TIME IF THEY ARE SPECIFIED WITHIN THE SAME T-STEP AS THE XIO SIGNAL. THE FOLLOWING PROGRAM SEGMENT PLACES THE VALUE "4" ON THE COMMAND LINES, THE VALUE "75" ON THE DEVICE SELECT LINES, AND SENDS XIO TO PORT 3.

```
.... KA = 3  "PORT NUMBER", KB = 75, KX = 4
X... G(6 KB)  "VALUE OF KB (=75) OUTPUT ON DEVICE SELECT LINES"
      KX->FIPH "VALUE OF KX (=04) OUTPUT ON CPU COMMAND LINES"
      XIO     "XIO IS SENT TO PORT 3"
```

SECTION 4.6 DESCRIBES A NUMBER OF SIGNALS USED FOR TIMING AND INTERRUPT PURPOSES AS WELL AS A SPECIAL SIGNAL CALLED "MASTER CLEAR". THE MASTER CLEAR SIGNAL IS GENERATED WHEN THE CPU IS POWERED ON, OR WHEN THE "SYSTEM RESET" BUTTON IS DEPRESSED. THE SIGNAL IS AVAILABLE AT THE PORT SO ANY EXTERNAL DEVICES MAY BE INITIALIZED ALONG WITH THE CPU.

0.2.3 QM-1 INTERRUPT OPERATION

THIS SECTION DISCUSSES THE QM-1 INTERRUPT SYSTEM IN TERMS OF PROGRAM TIMING AND SEQUENCES RATHER THAN THE ACTUAL HARDWARE UNITS INVOLVED, HENCE A THOROUGH UNDERSTANDING OF SECTIONS 4.5.2 AND 5.8.2 IS NECESSARY BEFORE GOING FURTHER.

THE STARTING POINT FOR AN EXTERNAL QM-1 INTERRUPT IS AN INTERRUPT FROM SOME EXTERNAL SOURCE. THE INTERRUPT IS NOW IN THE LATCHED STATE. THERE IS NO WAY, WITHIN THE CPU, TO PREVENT AN INTERRUPT FROM LATCHING. ONCE LATCHED, THE INTERRUPT REMAINS LATCHED UNTIL ACKNOWLEDGED IN THE NORMAL MANNER OR CLEARED BY THE GENERATE/CLEAR INTERRUPT MECHANISM (SEE SECTION 5.8.1) OR MASTER CLEAR. THE LATCHED INTERRUPT MUST BE ENABLED BY THE PROGRAMMER IN ORDER TO BECOME PENDING. A LATCHED INTERRUPT BECOMES PENDING FOLLOWING THE FIRST GATE NS PRIMITIVE AFTER THE CORRESPONDING ENABLE BIT IS SET. SIMILARLY, AN ENABLED INTERRUPT BECOMES PENDING ON THE FIRST GATE NS AFTER IT IS LATCHED. IF AN INTERRUPT IS PENDING WHEN ITS CORRESPONDING ENABLE BIT IS TURNED OFF THE INTERRUPT REMAINS PENDING UNTIL THE NEXT GATE NS. THIS MEANS THAT THE NANO-WORD WHICH CLEARS ANY ENABLE BITS SHOULD NOT ALLOW INTERRUPTS!

ONCE AN INTERRUPT IS PENDING IT BECOMES AVAILABLE FOR PRIORITY SELECTION. AS MENTIONED IN SECTION 4.5.2, THE PRIORITY SELECTION MECHANISMS ARE ACTIVATED/DEACTIVATED BY THE ALLOW NANO AND ALLOW MICRO INTERRUPT NANOPRIMITIVES. A TIMING CONFLICT ARISES WHEN READ NS (OF NPC) OCCURS IN T1 AND EITHER ALLOW INTERRUPT CONTROL IS ALSO SPECIFIED. THE RESULT OF THE READ NS IS UNDEFINED SINCE THE READ IS LEADING EDGE AND THE PRIORITY MECHANISM WAS ACTIVATED ON THE SAME LEADING EDGE. THE OBVIOUS RULE OF THUMB IS NOT TO READ NS (OF NPC) IN T1 OF ANY NANOWORD WHICH ALSO ALLOWS INTERRUPTS.

THE OVERALL TIMING OF AN INTERRUPT, FROM LATCHING TO ACKNOWLEDGEMENT, IS A FUNCTION OF THE EXECUTING NANOCODE. THE THREE OPERATIONS NECESSARY TO HAVE A LATCHED INTERRUPT ACKNOWLEDGED (ASSUMING IT IS ENABLED AND ALLOWED) ARE:

1. GATE NS CAUSES THE INTERRUPT TO BECOME PENDING.
2. READ NS (OF NPC) READS THE INTERRUPT SELECTED NANOCODE.
3. GATE NS GATES THE INTERRUPT NANOCODE AND CLEARS THE INTERRUPT LATCH AND PENDING FLAGS.

AN INTERRUPT MAY ALSO BE CLEARED BY GENERATE/CLEAR INTERRUPT (SECTION 5.8.1). THE CLEAR INTERRUPT FUNCTION CLEARS BOTH THE LATCHED AND PENDING FLAGS OF AN INTERRUPT. THE FUNCTION IS COMPLETED BY THE TRAILING EDGE OF T1.

8.3 STANDARD CHANNEL CONTROLLER

8.3.1 GENERAL

THE BEHAVIOR OF THE STANDARD QM-1 I/O CHANNEL IS SIMILAR TO THAT KNOWN THROUGHOUT THE COMPUTER INDUSTRY AS A "MULTIPLEXOR" CHANNEL. THIS MEANS THAT SEVERAL INDEPENDENT I/O DEVICES MAY BE CONCURRENTLY TRANSFERRING DATA OVER THE SAME I/O CHANNEL, WITHOUT KNOWLEDGE OF EACH OTHERS EXISTENCE. OF MAJOR INTEREST AT THIS POINT IS THE FACT THAT UNLIKE A CONVENTIONAL MULTIPLEXOR CHANNEL, WHICH SUPPORTS EITHER LOW SPEED DEVICES ONLY OR HIGH SPEED DEVICES WHEN THE LOW SPEED DEVICES ARE INACTIVE, THE QM-1 STANDARD CHANNEL WILL SUPPORT AN ACTIVE MIXTURE OF BOTH CLASSES OF DEVICES. THIS IS FEASIBLE AS LONG AS THEIR AGGREGATE DATA RATES DO NOT EXCEED THE ACCESS SPEED OF THE DESTINATION MEMORY. FOR EXAMPLE, A HIGH SPEED DISK TRANSFER OF 800,000 BYTES PER SECOND, A TAPE DRIVE TRANSFER OF 90,000 BYTES PER SECOND, AND SEVERAL LOW SPEED UNIT RECORD DEVICES (IE. PRINTERS, PUNCHES) COULD BE ACTIVE TOGETHER OVER THE SAME DMA PATH TO MAIN STORE WITHOUT ENCOUNTERING LOSS OF DATA DUE TO INTERFERENCE WITHIN THE CHANNEL. OF COURSE THE SYSTEM DESIGNER SHOULD BE AWARE THAT THERE IS ALWAYS THE HAZARD OF LOST DATA DUE TO ONE OR MORE OF THE DESIRED MEMORY BANKS BEING TIED UP THROUGH ACTIVITY FROM OTHER STORAGE PORTS. THE ABOVE MIXTURE WARRANTS BEING ASSIGNED TO THE HIGHEST PRIORITY MAIN STORE PORT.

THE QM-1 CHANNEL CONTROLLER COORDINATES THE COMMUNICATION BETWEEN THE QM-1 CPU AND I/O DEVICE CONTROLLERS, AND SYNCHRONIZES THE DEMANDS FOR DATA AND STATUS INTERRUPTS TO THE CPU. THE RELATIONSHIP OF THE CHANNEL CONTROLLER TO THE OTHER ELEMENTS OF THE I/O SYSTEM IS SHOWN IN FIGURE 8.3.1A. WHENEVER THE CPU SENDS AN XIO SIGNAL THE CONTROLLER WILL EITHER RESPOND TO THE COMMAND OR WILL IMMEDIATELY PASS THE COMMAND TO THE DEVICE CONTROLLER INDICATED. I/O COMMANDS ARE DIVIDED INTO TWO CATEGORIES, CHANNEL COMMANDS AND DEVICE COMMANDS. EACH COMMAND IS RECOGNIZED AS A SIX BIT QUANTITY, FOUND ON THE CPU COMMAND LINES DURING AN XIO SIGNAL (COMMAND LINES ORIGINATE ON THE PHANTOM BUS, SEE SECTION 8.2.2). COMMANDS WITH VALUES IN THE RANGE OF 00 - 67 (OCTAL) ARE CONSIDERED DEVICE COMMANDS AND ARE PASSED DIRECTLY TO THE DEVICE CONTROLLER IDENTIFIED BY THE SIX BIT QUANTITY ON THE DEVICE SELECT LINES (ORIGINATING FROM THE G BUS, SEE SECTION 8.2.2). COMMANDS WITH VALUES 70 - 77 (OCTAL) ARE CHANNEL COMMANDS AND RESULT IN IMMEDIATE CHANNEL CONTROLLER ACTION. THE SPECIFIC CHANNEL COMMANDS ARE DISCUSSED IN SECTION 8.3.2.

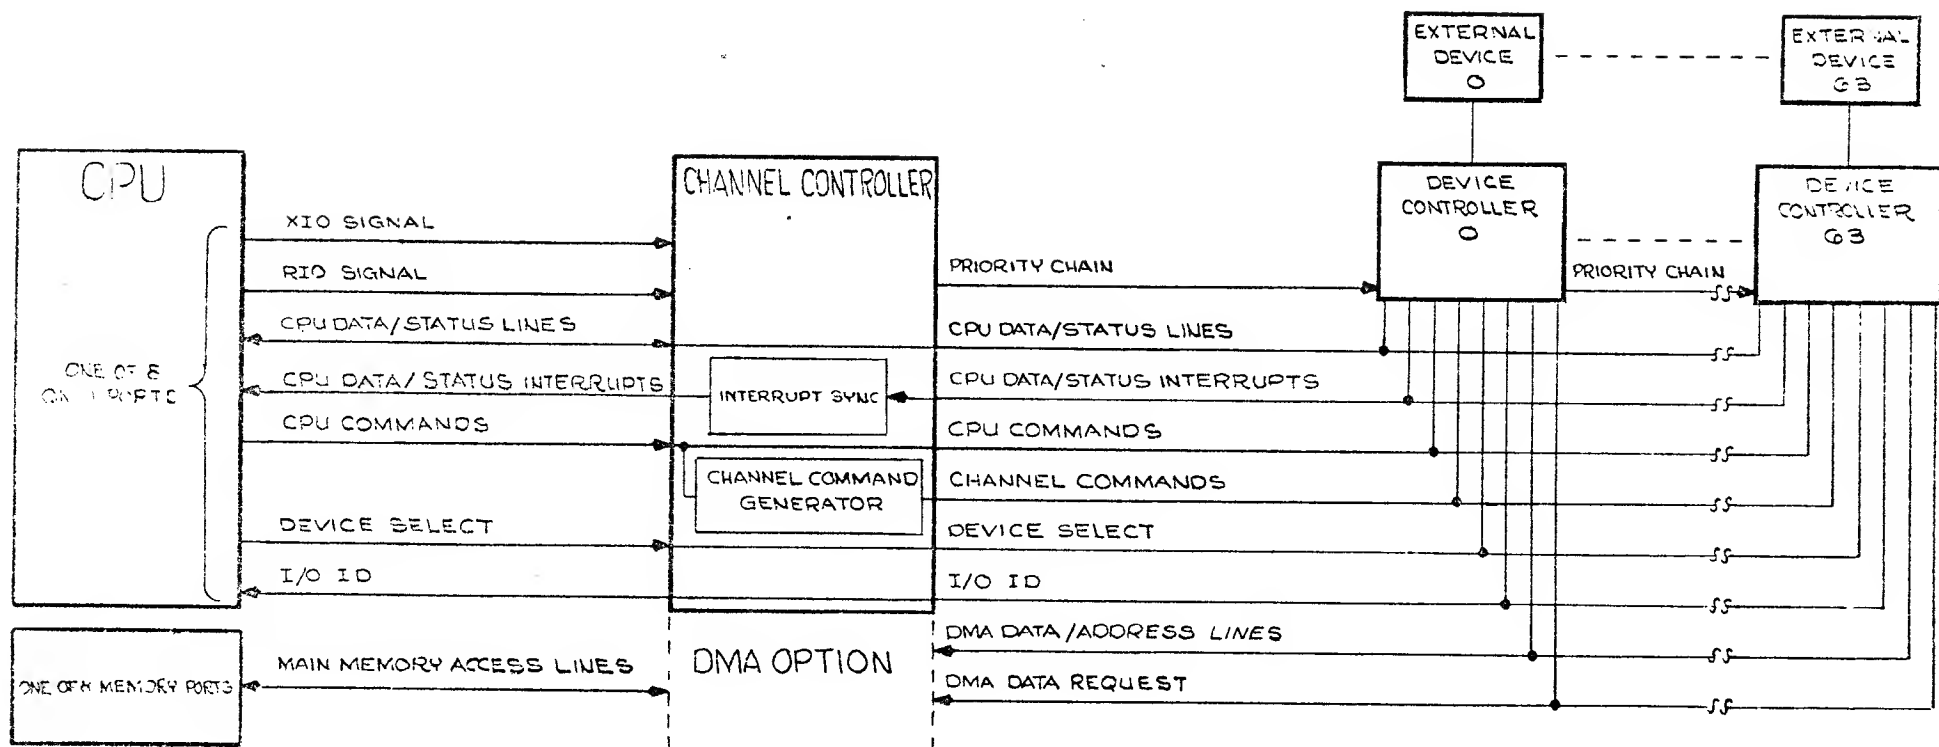


FIG 8.3.1A - CHANNEL CONTROLLER

THE CHANNEL MAY BE CONSIDERED TO HAVE THREE MODES OF OPERATION: DIRECT REQUEST, DATA, AND STATUS. THE ACTUAL MODE OF OPERATION IS MAINTAINED JOINTLY BY THE CHANNEL CONTROLLER AND EACH DEVICE CONTROLLER (DISCUSSED IN SECTION 8.4). DIRECT REQUEST MODE IS A MOMENTARY CHANNEL STATE WHICH EXISTS BETWEEN THE TIME OF ISSUANCE OF A DIRECT REQUEST COMMAND (VIA X10) AND THE GATING OF THE DESIRED DATUM (VIA RIO). AN EXAMPLE OF AN ACTUAL DIRECT REQUEST COMMAND IS "READ WORD COUNT" REGISTER. THIS COMMAND IS TRANSMITTED TO A SPECIFIC DEVICE CONTROLLER. THE CHANNEL CONTROLLER LOCKS ITSELF INTO THIS OPERATION UNTIL AN RIO SIGNAL IS RECEIVED FROM THE CPU, GATING THE REQUESTED INFORMATION INTO THE APPROPRIATE I/O PORT (ES) REGISTER. IT IS THE RESPONSIBILITY OF THE REQUESTING CPU PROGRAM TO ALLOW FOR THE PROPAGATION TIME OF THE ORIGINAL COMMAND TO REACH THE SPECIFIED DEVICE CONTROLLER AND FOR THE DESIRED INFORMATION TO RETURN TO THE PORT DATA LINES. FOR STANDARD I/O CABLE LENGTHS (UP TO 75 FEET FROM THE CPU TO THE CHANNEL AND DEVICE CONTROLLERS) A 300 NANOSECOND ROUND TRIP IS REQUIRED. THE FOLLOWING IS AN EXAMPLE OF A DIRECT REQUEST PROCEDURE, FOLLOWED BY A TIMING DIAGRAM


```

GET STATUS:          "FROM CHANNEL A, DEVICE B"
....   KA = ** "PORT SELECTION", KB = RD.WORDCT "STATUS COMMAND VALUE"
....   HOLD 2  "RETAIN KA AND KB", BRANCH( READ PORT ) "NEXT NANOWORD"
S...   A->KA        "SET PORT NUMBER INTO PORT SELECT (KA)"
.X...           "PORT SELECTION SETTling TIME"
...S... G(G B)      "DEVICE SELECTION NUMBER FROM B PARAMETER OF R31"
...S... KB->FIPH     "COMMAND VALUE TO COMMAND LINES"
...S... X10         "INITIATE COMMAND TRANSMISSION"
          "AT LEAST 5 T-PERIODS SHOULD ELAPSE FROM LEADING EDGE OF X10"
...S... READ NS,   GATE NS

READ PORT:          "PORT NUMBER PASSED VIA KA, USING HOLD 2"
....
X...   READ NS
.X...   RIO        "PLACE DESIRED COUNTER INTO PORT REGISTER"
.X...   GATE NS    "EXIT THIS PROCEDURE"

```

** INDICATES AN UNDEFINED VALUE, TO BE SET UP DURING PROGRAM EXECUTION.



WHILE DIRECT REQUEST MODE IS INITIATED BY THE CPU, BOTH DATA AND STATUS MODES ARE SELECTED FROM DEMANDS OF THE DEVICE CONTROLLERS. THE ACTUAL NOTIFICATION AT THE CPU, OF CHANNEL ENTRY INTO ONE OF THESE MODES IS THROUGH DATA OR STATUS INTERRUPTS. A STATUS INTERRUPT ORIGINATES AT AN I/O DEVICE, AND INDICATES EITHER A CHANGE OF STATE WITHIN THE DEVICE OR AN END OF OPERATION CONDITION. ALL CONDITIONS CAPABLE OF TRIGGERING STATUS INTERRUPTS MAY BE MASKED, UNDER PROGRAMMED CONTROL, AS TO WHETHER THEIR OCCURRENCE MAY REQUEST THE INTERRUPT OR ONLY SET A STATUS INDICATOR. THE CHANNEL CONTROLLER IS CONTINUALLY INTERROGATING ITS DEVICE CONTROLLERS FOR THE PRESENCE OF A STATUS INTERRUPT REQUEST FLAG. THIS INTERROGATION PROCEEDS FROM THE CHANNEL CONTROLLER, DOWN THE PRIORITY CHAIN LINE, THROUGH EACH DEVICE CONTROLLER, IN THE ORDER OF THEIR CONNECTION ON THE PRIORITY CHAIN. ONCE A STATUS INTERRUPT REQUEST FLAG IS RECOGNIZED THE CHANNEL CONTROLLER LOCKS OUT FURTHER STATUS REQUEST ANALYSIS AND TRANSMITS A STATUS INTERRUPT TO THE CPU. AT THIS TIME THE DEVICE IDENTIFICATION (IO ID) OF THE INTERRUPTING DEVICE, AND ITS INTERRUPT STATUS REGISTER, IS RETAINED UNTIL CALLED FOR BY THE CPU. NO FURTHER STATUS REQUESTS WILL BE ACCEPTED BY THE CHANNEL CONTROLLER UNTIL THE CPU HAS COMPLETED THE HANDLING OF THIS INTERRUPT.

THE ISSUANCE OF A STATUS INTERRUPT TO THE CPU PREPARES THE CHANNEL TO ENTER STATUS MODE. THIS MODE IS NOT ACTUALLY ENTERED UNTIL THE CPU PROCESSES THE STATUS INTERRUPT. STATUS MODE MAY BE CONSIDERED A MOMENTARY CHANNEL STATE, AS DESCRIBED FOR DIRECT REQUEST MODE ABOVE. IT IS ENTERED FOLLOWING ISSUANCE OF THE CHANNEL COMMAND "STATUS REQUEST" (DESCRIBED IN SECTION 8.3.2) AND IS TERMINATED WITH THE NEXT RIO SIGNAL. TIMING IS THE SAME AS FOR DIRECT REQUEST (ABOVE). THE STATUS INTERRUPT, "STATUS REQUEST", SEQUENCE IS THE ONLY METHOD PROVIDED FOR ACCESSING AN INTERRUPT STATUS REGISTER (SEE SECTION 8.4.2.2).

DATA INTERRUPT PROCESSING IS HANDLED INDEPENDENTLY OF STATUS INTERRUPTS. USING THE SAME PRIORITY ORDERED INTERROGATION PROCEDURE DEFINED FOR STATUS INTERRUPT REQUESTS (ABOVE), THE CHANNEL CONTROLLER SEARCHES FOR THE OCCURRENCE OF A DATA INTERRUPT REQUEST. THIS INTERROGATION IS NOT AFFECTED BY THE PRESENCE OF A PENDING STATUS INTERRUPT, NOR WILL THE INTERROGATION FOR STATUS REQUESTS BE AFFECTED BY A PENDING DATA INTERRUPT. AS SOON AS A DATA INTERRUPT REQUEST IS RECOGNIZED THE CHANNEL IS PLACED IMMEDIATELY INTO DATA MODE. THE SEARCH FOR ANY OTHER DATA INTERRUPT REQUESTS IS INHIBITED UNTIL THE CPU HAS PROCESSED THIS DATA INTERRUPT. THE CHANNEL REMAINS IN DATA MODE UNTIL THE ACTIVE DATUM IS PASSED BETWEEN THE CPU AND I/O DEVICE. PRIOR TO THE ACTUAL PROCESSING OF A DATA INTERRUPT, DATA MODE MAY BE TEMPORARILY SUPPRESSED BY CPU ACTION TO PLACE THE CHANNEL INTO EITHER DIRECT REQUEST OR STATUS MODE. WHEN THE MOMENTARY CHANNEL OPERATION IS COMPLETED THE CHANNEL WILL RETURN TO ITS PREVIOUS DATA STATE.

DATA INTERRUPT PROCESSING VARIES, DEPENDING ON THE DIRECTION OF DATA FLOW REQUESTED. A DEVICE CONTROLLER PROCESSING AN OUTPUT OPERATION WILL MAKE DATA-OUT INTERRUPT REQUESTS OF THE CHANNEL, WHILE INPUT OPERATIONS WILL MAKE DATA-IN INTERRUPT REQUESTS. WHEN THE CHANNEL CONTROLLER IS PROCESSING A DATA-IN INTERRUPT IT WILL REMAIN IN DATA MODE UNTIL EXACTLY TWO (2) RIO SIGNALS HAVE BEEN SENT BY THE CPU. THE FIRST RIO INDICATES THAT THE CPU IS BEGINNING TO HANDLE THE CURRENT INTERRUPT. THE INFORMATION PRESENT ON THE PORT DATA LINES, SINCE ENTERING DATA MODE, IS THE MEMORY ADDRESS AT WHICH THE FOLLOWING DATA WORD IS TO BE STORED. THIS ADDRESS ORIGINATES FROM A REGISTER WITHIN THE DEVICE CONTROLLER INVOLVED IN THIS DATA MODE OPERATION, AND IS DESCRIBED IN SECTION 8.4.3 (DATA ROUTING SUPPORT). THE RIO GATES THE ADDRESS INTO THE PORT REGISTER AND THEN ALLOWS THE CHANNEL CONTROLLER TO FETCH THE ACTUAL DATA WORD FROM THE DEVICE CONTROLLER. THE SECOND RIO GATES THE DATA WORD INTO THE PORT REGISTER, AND RELEASES THE CHANNEL CONTROLLER FROM DATA MODE. INTERROGATION FOR THE NEXT DATA INTERRUPT REQUEST CAN NOW BEGIN. THE CPU PROGRAM HANDLING THE INPUT DATA OPERATION MUST ALLOW FOR PROPAGATION TIME BETWEEN THE TWO RIO SIGNALS. IT TAKES APPROXIMATELY 300 NANSECONDS FOR THE FIRST SIGNAL TO REACH THE DEVICE CONTROLLER AND FOR THE ACTUAL DATA TO RETURN TO THE PORT DATA LINES (75 FOOT CABLE LENGTHS).

DATA-OUT INTERRUPT PROCESSING DIFFERS SLIGHTLY FROM DATA-IN OPERATION. EXACTLY TWO SIGNALS ARE REQUIRED TO RELEASE THE CHANNEL FROM DATA MODE. IN THIS CASE AN RIO SIGNAL WILL GATE A MEMORY ADDRESS INTO THE PORT REGISTER, AS WITH DATA-IN ABOVE, AND AS SOON AS THE SPECIFIED DATA WORD HAS BEEN PLACED INTO THE PORT REGISTER THE CHANNEL COMMAND "DATA AVAILABLE" (DESCRIBED IN SECTION 8.3.2) IS ISSUED WITH AN XIO SIGNAL. THIS COMMAND TRANSMITS THE DATUM TO THE I/O DEVICE AND RELEASES DATA MODE.

DURING STATUS MODE AND DATA MODE OPERATIONS THE CPU NEED NOT PROVIDE DEVICE SELECTION INFORMATION ON THE G-BUS LINES. THIS IS DUE TO THE CHANNEL CONTROLLER UNCONDITIONALLY LOCKING INTO THE DEVICE CONTROLLER ASSOCIATED WITH THE INTERRUPT. ON THE OTHER HAND, THE CPU USUALLY DOES NEED TO KNOW THE IDENTITY OF THE DEVICE CAUSING AN INTERRUPT. FOR THIS REASON EACH DEVICE CONTROLLER WILL PRESENT ITS UNIQUE DEVICE SELECTION NUMBER TO THE CHANNEL'S "IO ID" LINES DURING STATUS MODE. THE CORRECT "IO ID" WILL BE AVAILABLE 300 NANOSECONDS AFTER THE ISSUANCE OF THE "STATUS REQUEST" COMMAND, AND WILL REMAIN STABLE UNTIL AFTER THE RIO RELEASE SIGNAL. DURING DATA MODE INTERRUPTS THE DEVICE ID IS NORMALLY UNNECESSARY, SINCE THE PRESENCE OF THE MEMORY ADDRESS IS ALL THAT IS NEEDED TO COMPLETE THE DATA TRANSFER CYCLE. WHEN USING QM-1 VIRTUAL MEMORY OPTIONS AN ADDITIONAL IDENTIFIER, CALLED THE "TASK ID", IS REQUIRED TO SELECT THE APPROPRIATE STORAGE PARTITION. THE "TASK ID" REPLACES THE DEVICE ID ON THE "IO ID" LINES, AND IS ORIGINALLY PRESET BY THE OPERATING SYSTEM UNDER PROGRAMMED CONTROL (SEE SECTION 8.4.2.1, DEVICE CONTROL WORD).

8.3.2 CHANNEL COMMANDS

SEVEN OF THE EIGHT POSSIBLE CHANNEL COMMANDS ARE CURRENTLY ASSIGNED. THEIR FUNCTION AND USAGE ARE DESCRIBED BELOW. COMMAND CODE VALUES ARE SHOWN IN OCTAL. TRANSMISSION OF A CHANNEL COMMAND REQUIRES PLACEMENT OF THE COMMAND VALUE ON THE PHANTOM BUS LINES DURING ISSUANCE OF THE XIO SIGNAL. THE COMMAND IS THEN SENT TO THE CHANNEL CONTROLLER ON THE "KA" SELECTED I/O PORT.

- I/O RESET (70) - RESETS THE CHANNEL CONTROLLER AND ALL DEVICE CONTROLLERS ON THAT CHANNEL. EACH CONTROLLER IS SET TO ITS INITIAL STATE. ALL INTERNAL REGISTERS ARE CLEARED TO ZERO, AND ALL INTERRUPT FLAGS AND STATUS INDICATORS ARE RESET. THIS COMMAND SHOULD BE USED ONLY DURING SYSTEM INITIALIZATION.
- DATA AVAILABLE (71) - USED ONLY DURING DATA-OUT MODE CHANNEL OPERATION, FOLLOWING THE MEMORY ADDRESS GATE RIO SIGNAL. DATA AVAILABLE IS ISSUED FOLLOWING PLACEMENT OF THE ACTUAL DATA INTO THE APPROPRIATE PORT REGISTER. DATA-OUT MODE IS CLEARED IMMEDIATELY FOLLOWING THIS COMMAND.
- (72) - UNASSIGNED AT PRESENT.
- ENABLE DATA INTERRUPTS (73) - ALLOWS DATA INTERRUPT REQUESTS TO BE ACCEPTED BY THE CHANNEL CONTROLLER AND TO BE ISSUED TO THE CPU. SINCE INITIAL CHANNEL CONDITIONS INHIBIT INTERRUPTS THIS COMMAND MUST BE USED PRIOR TO DATA TRANSFER OPERATIONS; FOLLOWING A "DISABLE DATA INTERRUPTS" COMMAND (74), AN "I/O RESET" COMMAND (70), OR SYSTEM "MASTER CLEAR" MANUAL OPERATION.
- DISABLE DATA INTERRUPTS (74) - BLOCKS THE SELECTION OF DATA INTERRUPTS BY THE CHANNEL CONTROLLER. THIS IS AN INITIAL STATE CONDITION (I/O RESET), AND MAY BE CLEARED WITH AN "ENABLE DATA INTERRUPTS" COMMAND (73).
- ENABLE STATUS INTERRUPTS (75) - ALLOWS STATUS INTERRUPT REQUESTS TO BE ACCEPTED BY THE CHANNEL CONTROLLER AND TO BE ISSUED TO THE CPU. SINCE INITIAL CHANNEL CONDITIONS INHIBIT INTERRUPTS THIS COMMAND MUST BE USED PRIOR TO NORMAL CHANNEL ACTIVITY; FOLLOWING A "DISABLE STATUS INTERRUPTS" COMMAND (76), AN "I/O RESET" COMMAND (70), OR A SYSTEM "MASTER CLEAR" MANUAL OPERATION.

DISABLE STATUS INTERRUPTS (76) - BLOCKS THE SELECTION OF STATUS INTERRUPTS BY THE CHANNEL CONTROLLER. THIS IS AN INITIAL STATE CONDITION (I/O RESET), AND MAY BE CLEARED WITH AN "ENABLE STATUS INTERRUPTS" COMMAND (75).

STATUS REQUEST (77) - USED ONLY FOLLOWING A STATUS INTERRUPT, IN ORDER TO PLACE THE CHANNEL INTO STATUS MODE. FOLLOWING ISSUANCE OF THIS COMMAND THE INTERRUPTING DEVICE CONTROLLER PLACES ITS INTERRUPT STATUS REGISTER ONTO THE PORT DATA LINES (UP TO 18 BITS OF INFORMATION) AND ITS DEVICE IDENTIFICATION ON THE "IO ID" LINES (6 BITS). FOLLOWING APPROPRIATE COMMAND/DATA PROPAGATION DELAY (USUALLY 300 NANoseconds) AN RIO SIGNAL WILL GATE THE INTERRUPT STATUS REGISTER DATA INTO THE I/O PORT REGISTER AND RELEASE THE CHANNEL FROM STATUS MODE.

8.4 STANDARD DEVICE CONTROLLER

8.4.1 GENERAL

THE STANDARD DEVICE CONTROLLER MAINTAINS CONTROL OVER ONE OR MORE PHYSICAL I/O UNITS, AND ACTS AS THE INTERFACE BETWEEN THESE UNITS AND THE STANDARD CHANNEL CONTROLLER. THE DEVICE CONTROLLER IS DESIGNED TO ALLOW GENERALIZED SOFTWARE TO HANDLE A VAST ARRAY OF VERY DIFFERENT I/O EQUIPMENT. IT CONTAINS A NUMBER OF FEATURES WHICH MAY, OR MAY NOT, BE USED BY THE SOFTWARE, ALLOWING A VARYING DEGREE OF HARDWARE SUPPORT. FIGURE 8.4.1A SHOWS THE RELATIONSHIP OF THE DEVICE CONTROLLER TO THE STANDARD CPU, DIRECT MEMORY ACCESS CHANNELS AND THE EXTERNAL I/O DEVICES. THIS FIGURE ALSO SHOWS THE MAJOR HARDWARE FEATURES OF THE DEVICE CONTROLLER. THESE FEATURES ARE DIVIDED INTO TWO MAJOR CATEGORIES: DEVICE CONTROL SUPPORT AND DATA TRANSFER SUPPORT.

WITHIN THE DEVICE CONTROLLER SOME OF THE CONTROL SUPPORT UNITS CAN BE CLASSIFIED AS "STATIC CONTROLS". THESE ARE THE "DEVICE CONTROL WORDS", THE "STATUS REGISTER", AND THE "INTERRUPT STATUS REGISTER". THESE ARE REFERRED TO AS "STATIC" CONTROLS SINCE THEY ALL CONTAIN RESIDUAL INFORMATION TO BE USED BY THE REST OF THE I/O SYSTEM, USUALLY BEYOND THE SCOPE OF ONE DEVICE OPERATION.

THE REMAINING CONTROL SUPPORT UNITS ARE CLASSIFIED "DYNAMIC CONTROLS". THESE ARE THE UNITS WHICH RESPOND DIRECTLY TO COMMANDS FROM THE CPU, THE CHANNEL CONTROLLER OR THE DEVICE ITSELF.

ADDITIONAL UNITS ARE PROVIDED FOR DATA ROUTING SUPPORT. THESE ARE THE BUFFER ADDRESS REGISTER (BAR), WHICH SERVES AS A MEMORY ADDRESS POINTER, AND THE WORD COUNT REGISTER (WC), WHICH CAN BE USED TO COUNT THE NUMBER OF WORDS, OR CHARACTERS, TO BE TRANSFERRED. WC CAN BE DECREMENTED AND BAR CAN BE INCREMENTED OR DECREMENTED AUTOMATICALLY DURING DATA TRANSFERS. THESE REGISTERS MAY BE APPLIED TO DATA TRANSFERS OVER THE DATA PATH TO THE ACTUAL QM-1 PORT REGISTER, OR DIRECTLY TO MAIN STORE WHEN THE DIRECT MEMORY ACCESS OPTION IS INSTALLED IN THE CHANNEL CONTROLLER. FINALLY, A "DATA CHAINING" FACILITY IS PROVIDED WHICH ALLOWS THE SOFTWARE TO COMBINE DISJOINT BLOCKS OF MEMORY INTO ONE LOGICALLY CONTINUOUS I/O BUFFER, WITH MINIMAL SOFTWARE INTERVENTION. THE DEVICE CONTROLLER ALSO CONTAINS THE HARDWARE NECESSARY TO RECOGNIZE AND PROCESS INTERRUPT REQUESTS FROM THE DEVICE, AS WELL AS TO GENERATE SPECIAL INTERRUPTS ON DEVICE CONTROLLER STATUS OR DMA CHANNEL STATUS CONDITIONS.

8.4.2 DEVICE CONTROL

8.4.2.1 DEVICE CONTROL WORDS

THE DEVICE CONTROL WORDS CONSIST OF TWO REGISTERS CONTAINING "STATIC" CONTROL INFORMATION, SUCH AS "SELECT DMA OPERATING MODE", "ACTIVATE DATA TRANSLATION", ETC. THE TWO REGISTERS ARE REFERRED TO AS DCWA AND DCWB. THEIR CONTENT VARIES DEPENDING ON THE DEVICES ATTACHED TO THE CONTROLLER. THE DCW ALLOWS THE PROGRAMMER TO SELECT THE FUNCTIONAL UNITS OF THE DEVICE CONTROLLER THAT WILL BE UTILIZED IN AN OPERATION. FOR EXAMPLE: DCWA ALLOWS THE PROGRAMMER TO SELECT EVENTS ON WHICH STATUS INTERRUPTS WILL OCCUR.

DCWA AND DCWB ARE LOADED BY SPECIFIC DEVICE COMMANDS (SEE SECTIONS 8.4.2.5 AND 8.5.3), AND ARE ACCESSIBLE TO THE PROGRAMMER AT ANY TIME VIA THE DIRECT REQUEST MECHANISM (SECTIONS 8.3.1, 8.4.2.4, AND 8.5.6). NANODATA HAS ASSIGNED STANDARD (DEVICE INDEPENDANT) FUNCTIONS TO SEVERAL DCWA AND DCWB BIT POSITIONS, THAT ARE COMMON TO A WIDE RANGE OF EXTERNAL DEVICES. THE FOLLOWING LISTS DESCRIBE THEIR FUNCTIONS.

BIT IDENTITY	DESCRIPTION
DATA MAPPING CONTROLS	
DCWA 02 TRANSLATE	_CONTROLS CHARACTER CODE TRANSLATION WITHIN THE DEVICE CONTROLLER. THIS FEATURE IS OPTIONAL IN SOME DEVICES. EXAMPLE: TRANSLATE LOWER CASE ASCII TO UPPER CASE ASCII FOR PRINTING.
DCWA 03 PACK/UNPACK	_CONTROLS THE MAPPING OF CHARACTERS (BYTES) INTO THE STANDARD QM-1 18 BIT WORD. IN "PACKED" OPERATING MODE MOST DEVICE CONTROLLERS TREAT THE NC REGISTER AS A CHARACTER COUNTER, AND DECREMENT BY TWO FOR EACH QM-1 WORD TRANSFERRED.
STATUS INTERRUPT ENABLE MASKS	
DCWA 04 COMMAND REJECT	_SEE SECTIONS 8.4.2.2 AND 8.4.2.5 FOR COMPLETE INFORMATION ON "COMMAND REJECT".

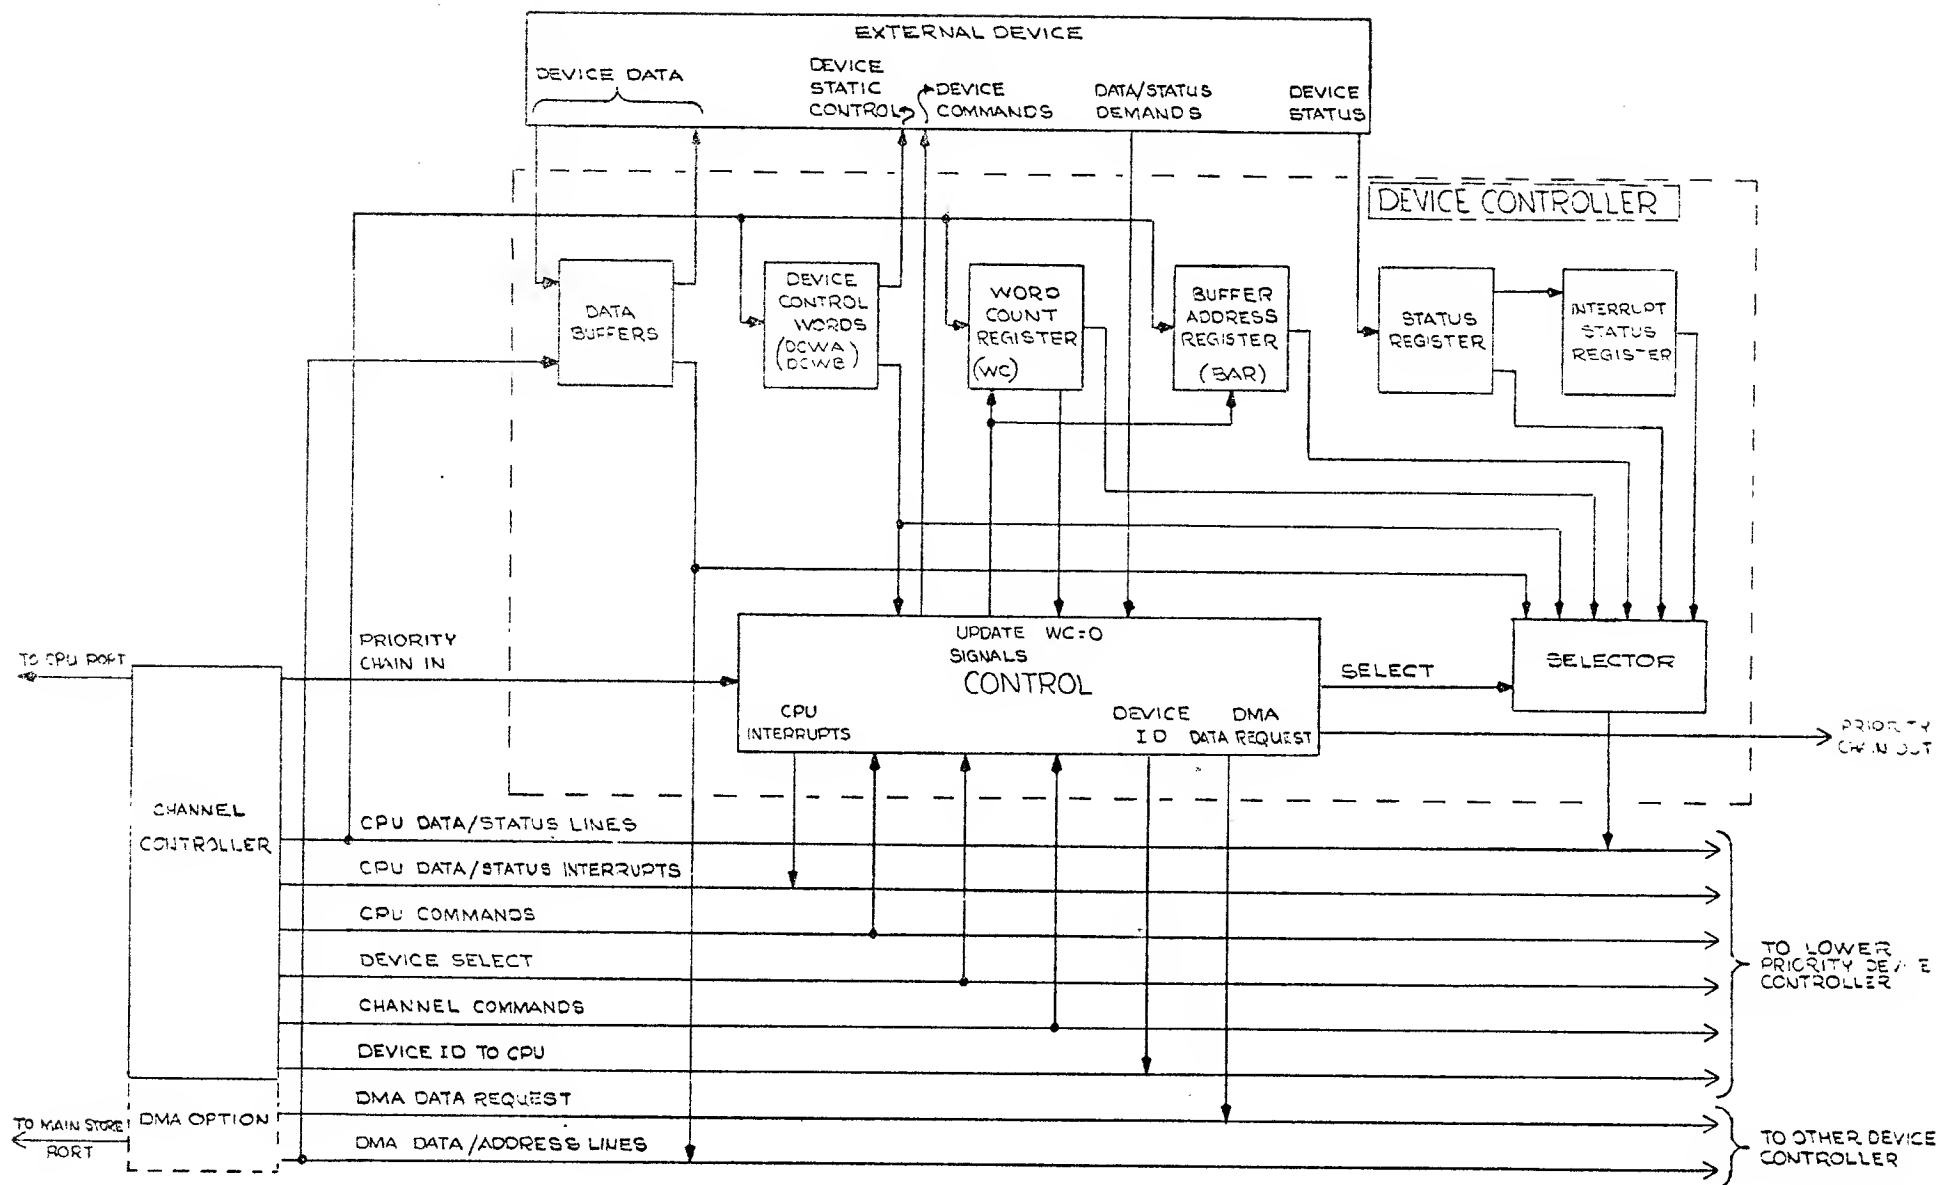


FIG 8.4.1A- DEVICE CONTROLLER

DCNA 05	ANY ERROR CONDITION	_SEE SECTION 8.4.2.2 FOR A DESCRIPTION OF THE "ERROR" STATUS BIT CONTROLS.
DCNA 06	DEVICE READY	_ALLOWS INTERRUPT ON DEVICE STATE CHANGING FROM READY TO NOT READY, OR NOT READY TO READY, CONDITION. SEE SECTION 8.4.2.2.
DCNA 07	NOT BUSY	_ALLOWS INTERRUPT ON DEVICE STATE CHANGING FROM BUSY (PERFORMING AN OPERATION) TO NOT BUSY (DEVICE INACTIVE). SEE SECTION 8.4.2.2 FOR ADDITIONAL DISCUSSION.
DCNA 08	UNIT AVAILABLE	_ALLOWS INTERRUPT TO INDICATE THAT THE IS READY TO RECEIVE COMMANDS FOR ITS NEXT OPERATING CYCLE. SEE SECTION 8.4.2.2.
DCNA 09	WORD COUNT ZERO	_ALLOWS INTERRUPT ON WORD COUNT REGISTER BEING DECREMENTED TO ZERO. TERMINATES CURRENT I/O TRANSFER. SEE 8.4.2.2.

IDENTIFICATION

DCNB 00	TASK ID 0	_TASK ID IS SIMPLY A SIX BIT FIELD, THE CONTENT OF WHICH IS PLACED ONTO THE DEVICE ID LINES DURING DATA INTERRUPTS, BY THE DEVICE CONTROLLER. ITS PURPOSE IS TO PROVIDE ADDITIONAL DATA ROUTING INFORMATION FOR THE SYSTEM UTILIZING THE DEVICE. SEE SECTION 8.4.3.3 FOR ADDITIONAL DETAILS.
DCNB 01	" " 1	
DCNB 02	" " 2	
DCNB 03	" " 3	
DCNB 04	" " 4	
DCNB 05	" " 5	

DATA ROUTING SUPPORT CONTROLS

DCNB 06	DECREMENT WORD COUNT (WC)	_PERMITS THE CONTENT OF THE WC REGISTER TO BE DECREMENTED FOLLOWING EACH DATA TRANSFER CYCLE. SEE SECTIONS 8.4.3.2 AND 8.5.4.
DCNB 07	INCREMENT BUFFER ADDRESS REGISTER (BAR)	_PERMITS THE CONTENT OF BAR TO BE INCREMENTED FOLLOWING EACH WORD TRANSFER CYCLE. SEE SECTIONS 8.4.3.1 AND 8.5.4.
DCNB 08	DECREMENT BUFFER ADDRESS REGISTER (BAR)	_SAME AS DCNB 07, EXCEPT DECREASES REGISTER VALUE. SEE SECTIONS 8.4.3.1 AND 8.5.4.
DCNB 09	DMA MODE	_SWITCHES THE DEVICE TO DMA OPERATING MODE. DMA OPTION IS REQUIRED. SEE 8.4.3.5.

NOTE: DCNA AND DCNB ARE CLEARED (ZEROED) BY MASTER CLEAR, I/O RESET, AND CLEAR DEVICE COMMANDS.

8.4.2.2 STATUS

EACH DEVICE CONTROLLER IS PROVIDED WITH TWO STATUS REGISTERS THAT INFORM THE I/O PROCESSING PROGRAM OF THE STATE OF THE DEVICE AND ITS ASSOCIATED CONTROLLER. THESE TWO REGISTERS ARE CALLED THE "STATUS REGISTER" (SR) AND THE "INTER-RUPT STATUS REGISTER" (ISR). THE STATUS REGISTER CONTAINS FULL AND DETAILED INFORMATION ON THE CONDITION OF THE DEVICE AND CONTROLLER. THE STATUS REGISTER IS ACCESSIBLE TO THE PROGRAM AT ANY TIME VIA THE "DIRECT REQUEST" MECHANISM (DESCRIBED IN SECTIONS 8.3.1 AND 8.4.2.4).

THE INTERRUPT STATUS REGISTER (ISR) DIFFERS FROM THE STATUS REGISTER (SR) IN THE NATURE OF ITS BIT CONTENT AND OPERATION. THE ISR CONTAINS A MORE SPARSE FORM OF SR INFORMATION BITS. ONE BIT OF ISR MAY REPRESENT A CLASS OF BITS IN SR. FOR EXAMPLE, THE GENERAL "ERROR" BIT OF ISR REPRESENTS ALL ERROR CONDITIONS OF THE DEVICE AND DEVICE CONTROLLER: LOST DATA, BAD PARITY, ETC. IN A COMPLICATED DEVICE, WHERE MORE THAN ONE STATUS REGISTER MAY BE REQUIRED, ISR IS USED AS A POINTER TO THE STATUS REGISTER THAT CONTAINS THE DETAILED INFORMATION ON THE REASON FOR THE STATUS INTERRUPT. THE PRIMARY PURPOSE OF THE ISR IS TO INFORM THE I/O PROCESS OF WHICH STATUS CONDITION(S) CAUSED THE STATUS INTERRUPT CURRENTLY BEING HANDLED (SUCH AS NORMAL END OF OPERATION, A DEVICE ERROR, ETC.).

THE SETTING OF AN ISR BIT ALWAYS OCCURS ALONG WITH A STATUS INTERRUPT. IN ORDER FOR AN ISR BIT TO SET IT MUST FIRST BE ENABLED BY THE PRIOR SETTING OF ITS CORRESPONDING DCW BIT (SEE SECTION 8.4.2.1, DEVICE CONTROL WORD). THE SETTING OF AN ISR BIT OCCURS ONLY DURING CHANGE OF THE RELATED STATE, WHILE THE CONDITION IS ENABLED. FOR EXAMPLE, THE DEVICE "NOT BUSY" BIT WILL SET, AND GENERATE A STATUS INTERRUPT, ONLY ON THE CHANGE OF DEVICE STATE FROM "BUSY" TO "NOT BUSY" WHILE DCWA BIT 07 (ENABLE INTERRUPT ON "NOT BUSY") IS SET. IF THE DEVICE IS ALREADY NOT BUSY WHEN, BIT 07 OF DCWA IS CHANGED TO ENABLE THE INTERRUPT (MADE A ONE), NO INTERRUPT WILL OCCUR. THE DEVICE MUST CHANGE STATE FROM "BUSY" TO "NOT BUSY", WHILE THE INTERRUPT IS ENABLED, TO TRIGGER THE INTERRUPT.

IN ORDER FOR THE ACTIVE I/O PROCESSING PROGRAM TO READ THE ISR IT MUST ISSUE A "STATUS REQUEST" COMMAND. IT MAY ISSUE THIS COMMAND ONLY AFTER THE RECEIPT OF A STATUS INTERRUPT. SINCE THE CHANNEL CONTROLLER "REMEMBERS" THE IDENTITY OF THE DEVICE CONTROLLER CURRENTLY REQUESTING THE STATUS INTERRUPT, THE "STATUS REQUEST" COMMAND WILL BE ROUTED DIRECTLY TO THAT DEVICE. THIS FUNCTION IS DESCRIBED IN MORE DETAIL IN THE FOLLOWING SECTION ON INTERRUPTS. UPON RECEIPT OF A "STATUS REQUEST" THE DEVICE CONTROLLER WILL PUT THE CONTENT OF ITS ISR, AND ITS DEVICE SELECTION NUMBER, ON THE CHANNEL DATA, AND DEVICE ID, LINES TO

THE CPU. A TIME DELAY OF 300 NANOSECONDS IS REQUIRED BETWEEN THE "STATUS REQUEST" COMMAND AND THE RIO THAT READS THE STATUS DATA INTO THE PORT REGISTER. THE DEVICE CONTROLLER DOES NOT ALLOW ANY ADDITIONAL ISR BITS TO SET FOLLOWING ITS RECEIPT OF THE "STATUS REQUEST", UNTIL COMPLETION OF STATUS INTERRUPT PROCESSING BY THE NEXT RIO. ALL ISR BITS SET, AND THEIR CORRESPONDING INTERRUPT REQUESTS, PRIOR TO A "STATUS REQUEST" WILL BE CONSIDERED AS FULLY ACKNOWLEDGED FOLLOWING THE RESPONSE TO JUST ONE STATUS INTERRUPT. THEREFORE SEVERAL ENABLED STATUS INTERRUPT CONDITIONS, OCCURRING WITHIN CLOSE PROXIMITY OF EACH OTHER, CAN BE PROCESSED AS A SINGLE STATUS INTERRUPT. ALL ISR BITS SET PRIOR TO THE INTERRUPT WILL BE CLEARED AUTOMATICALLY FOLLOWING THE TRANSFER OF THE ISR TO THE CPU.

THERE ARE THREE METHODS BY WHICH STATUS BITS MAY BE RESET (CLEARED OR ZEROED). THE FIRST TYPE OF STATUS INFORMATION MAY BE REFERRED TO AS "REAL TIME". THIS MEANS THAT THE DEVICE CONDITION ITSELF HAS CONTROL OVER THE STATE OF THE STATUS BIT. "DEVICE READY", FOR EXAMPLE, BELONGS TO THE "REAL TIME" CATEGORY, SINCE IT WILL CHANGE STATE EVERY TIME THE DEVICE READY CONDITION CHANGES. THE SECOND TYPE OF STATUS ARE THE "ERROR" CONDITION BITS. THEY ARE CLEARED THROUGH ISSUANCE OF THE "CLEAR ERROR" COMMAND, WHICH IS DESCRIBED IN THE SECTION ON DEVICE COMMANDS (8.4.2.5 BELOW). FINALLY, THE STATUS BITS WHICH ARE RELATED TO WORD COUNT GOING ZERO ARE CLEARED BY THE LOADING OF A NEW VALUE INTO THE WORD COUNT REGISTER. ALL STATUS BITS, EXCEPT FOR "REAL TIME", ARE CLEARED WITH "MASTER CLEAR", "I/O RESET" (SEE SECTION 8.3.2, CHANNEL COMMANDS), "CLEAR DEVICE" (8.4.2.5), AS WELL AS "CLEAR ERROR" (8.4.2.5).

THE POSITIONS OF MOST ISR BITS CORRESPOND TO THOSE BITS WITHIN SR HAVING THE SAME GENERAL MEANING. FOR THIS REASON ONLY ONE STATUS BIT DESCRIPTION LIST APPEARS BELOW. ALL STATUS INDICATORS DESCRIBED APPEAR IN THE SAME POSITIONS IN BOTH ISR AND SR. THE FOLLOWING LIST SPECIFIES STANDARD INTERRUPT STATUS, AND STATUS, REGISTER BITS ASSIGNED BY NANODATA TO COVER A WIDE RANGE OF EXTERNAL DEVICES. BIT POSITION NUMBERS ARE SHOWN IN DECIMAL. ALL ISR INDICATORS, DESCRIBED BELOW, WILL BE SET ONLY IF ENABLED IN THE CORRESPONDING DCN.

BIT POSITION & IDENTITY

DESCRIPTION OF STATUS

0 - DEVICE READY

-(REAL TIME CONTROLLED IN SR)

"ISR" STATUS INTERRUPT GENERATED, AND BIT IS SET, WHENEVER A CHANGE OCCURS IN THE DEVICE READY CONDITION.

"SR" A 1 INDICATES THAT THE DEVICE IS CURRENTLY "READY", 0 INDICATES "NOT READY".

- 1 - DEVICE NOT BUSY
- _(REAL TIME CONTROLLED IN SR)
"ISR" STATUS INTERRUPT GENERATED, AND BIT IS SET, WHEN THE DEVICE CHANGES STATE FROM "BUSY" TO "NOT BUSY".
"SR" A 1 INDICATES DEVICE CURRENTLY "NOT BUSY", 0 INDICATES DEVICE IN OPERATION.
- 2 - ERROR CONDITION
- _(ISR" STATUS INTERRUPT GENERATED, AND BIT IS SET, WHENEVER ANY ERROR CONDITION OCCURS IN THE DEVICE OR CONTROLLER. ERROR CONDITIONS SETTING THIS BIT INCLUDE ANY OF THOSE LISTED UNDER STATUS BITS 12 THROUGH 17, AS WELL AS ANY ADDITIONAL DEVICE DEPENDENT ERRORS IN BITS 6 THROUGH 11.
"SR" A 1 INDICATES UNCLEARED, OR UNRESOLVED, ERROR CONDITIONS. A 0 MEANS NO OUTSTANDING ERROR CONDITIONS.
- 3 - UNIT AVAILABLE
- _(REAL TIME CONTROLLED IN SR)
"ISR" STATUS INTERRUPT GENERATED, AND BIT IS SET, WHEN THE DEVICE BECOMES READY FOR A NEW DATA TRANSFER CYCLE. FOR EXAMPLE, THE DATA BUFFER OF THE DEVICE IS RELEASED FROM ITS PREVIOUS OPERATION.
"SR" A 1 INDICATES UNIT IS NOW AVAILABLE FOR THE START OF THE NEXT DATA TRANSFER OPERATION. A 0 MEANS THE DEVICE IS IN FULL OPERATION AND IS NOT AVAILABLE FOR ANY NEW COMMANDS
- 4 - WORD COUNT REACHED ZERO
- _(STATUS INTERRUPT GENERATED, AND BIT IS SET, WHEN THE WORD COUNT REGISTER REACHES ZERO. THIS BIT IS CLEARED BY LOADING A NEW VALUE INTO THE WORD COUNT REGISTER.
- 5 - DATA CHAINING
- _(STATUS INTERRUPT GENERATED, AND BIT IS SET, (IN PLACE OF BIT 4) WHEN THE WORD COUNT REGISTER REACHES ZERO IN DATA CHAINING MODE. THIS OCCURS ONLY IF BIT 16 OF THE PRECEDING WORD COUNT REGISTER LOAD WAS A 1. CLEARED WHEN LOADING A NEW VALUE INTO WC.
- 6 - 11 DEVICE DEPENDENT
- _(ERRORS AND CONDITION INDICATORS UNIQUE TO DIFFERENT DEVICES (SEE EXAMPLE IN SECTION 8.6). APPEARANCE OF THESE BITS WITHIN ISR AS WELL AS SR IS ALSO DEVICE DEPENDANT.

BITS 12 THROUGH 16 APPEAR, AS DESCRIBED, ONLY IN SR. THESE POSITIONS HAVE DEVICE DEPENDENT FUNCTIONS WITHIN ISR.

- | | |
|---------------------------|--|
| 12 - DMA ADDRESSING ERROR | _DMA REFERENCE ATTEMPTED OUTSIDE OF MEMORY SPACE INSTALLED ON SYSTEM. FOR VIRTUAL MAIN STORE SYSTEMS, DMA REFERENCE TO A NON-RESIDENT PAGE. SETS THE "ERROR" BIT (2). RESET BY THE "CLEAR ERROR" COMMAND. |
| 13 - DMA EXCEPTION | _DMA ENCOUNTERED A MAIN STORE PARITY ERROR, OR OTHER HARDWARE MALFUNCTION. SETS THE "ERROR" BIT (2). RESET BY THE "CLEAR ERROR" COMMAND. |
| 14 - ILLEGAL DATA | _AN ILLEGAL DATA CHARACTER OR PATTERN HAS BEEN RECEIVED BY THE DEVICE CONTROLLER FROM THE DEVICE. FOR EXAMPLE, A CARD READER MAY HAVE ENCOUNTERED A NON-EBCDIC HOLLERITH CODE WHEN OPERATING IN EBCDIC MODE. SETS THE "ERROR" BIT (2). RESET BY THE "CLEAR ERROR" COMMAND. |
| 15 - TRANSLATOR ERROR | _TRANSLATOR MALFUNCTION HAS OCCURRED IN A DEVICE CONTROLLER EQUIPPED WITH A DATA TRANSLATOR. SETS THE "ERROR" BIT (2). RESET BY THE "CLEAR ERROR" COMMAND. |
| 16 - LOST DATA | _LOST DATA OCCURED DUE TO LATENCY OF THE CPU OR OVERLOADING OF THE DMA PATH. SETS THE "ERROR" BIT (2). RESET BY THE "CLEAR ERROR" COMMAND. |
| 17 - COMMAND REJECT | _ILLEGAL COMMAND SENT TO THE DEVICE CONTROLLER. USUALLY AN UNRECOGNIZED COMMAND OR A COMMAND ISSUED TOO EARLY OR OUT OF SEQUENCE. SETS THE "ERROR" BIT (2). RESET BY THE "CLEAR ERROR" COMMAND. THIS ERROR INDICATOR APPEARS IN BOTH SR AND ISR. |

8.4.2.3 INTERRUPT MECHANISM

INTERRUPTS ARE SIGNALS WHICH NOTIFY AN ACTIVE CPU PROCESS THAT SOME CHANGE OF STATE HAS OCCURRED, USUALLY AT AN I/O DEVICE. THERE ARE FIVE PHASES IN THE INTERRUPT SEQUENCE. THESE PHASES ARE CALLED: REQUESTED, LATCHED, PENDING, ACKNOWLEDGED, AND RELEASED (THESE NAMES SHOULD NOT BE CONFUSED WITH THE CPU INTERRUPT TERMINOLOGY DESCRIBED IN SECTIONS 4.5.2 AND 5.8.2). THREE FLAGS ARE INVOLVED WITH THIS SEQUENCE, THESE ARE THE "REQUEST FLAG" AND THE "LATCH FLAG", LOCATED WITHIN THE DEVICE CONTROLLER, AND THE "PENDING FLAG", WHICH IS LOCATED WITHIN THE CHANNEL CONTROLLER. STATUS AND DATA INTERRUPTS MUST BE ENABLED AT THE DEVICE CONTROLLER BEFORE THIS SEQUENCE CAN START.

THE INTERRUPT BECOMES "REQUESTED" WHEN A DEVICE SIGNALS ITS DEVICE CONTROLLER OF A CONDITION THAT HAS BEEN ENABLED TO TRIGGER AN INTERRUPT. AN INTERRUPT MAY ALSO BE "REQUESTED" WHEN A CONDITION WITHIN THE DEVICE CONTROLLER ITSELF REQUIRES AN INTERRUPT. THIS ACTION IS ASYNCHRONOUS AND INDEPENDENT OF THE CONTROLLER STATE.

WHEN THE CHANNEL CONTROLLER INTERROGATES DEVICE CONTROLLERS FOR INTERRUPTS, IT SEARCHES FOR "REQUEST FLAGS". UPON ENCOUNTERING ONE IT SETS THE APPROPRIATE INTERRUPT "LATCH FLAG" (ONE EXISTS FOR EACH DATA-IN, DATA-OUT, AND STATUS INTERRUPT). THE INTERRUPT IS NOW IN THE "LATCHED" PHASE. AFTER THE INTERRUPT HAS BEEN LATCHED, THE "PENDING FLAG" IN THE CHANNEL CONTROLLER WILL BE SET. THE SETTING OF THE "PENDING FLAG" IS IMMEDIATELY FOLLOWED BY A SIGNAL TO THE ASSOCIATED CPU INTERRUPT LINE (SEE SECTION 8.3.1, FOR FURTHER DISCUSSION).

SINCE MORE THAN ONE INTERRUPT MAY BE "LATCHED" SIMULTANEOUSLY IT IS NECESSARY FOR THE CHANNEL CONTROLLER TO SELECT THE ONE WITH THE HIGHEST PRIORITY. IN THE TIME INTERVAL BETWEEN THE "LATCHED" AND "PENDING" STATES THE CHANNEL CONTROLLER HAS SEARCHED THE "PRIORITY CHAIN" CONNECTED DEVICE CONTROLLERS AND HAS SELECTED THE HIGHEST PRIORITY (CLOSEST CONNECTION TO THE CHANNEL CONTROLLER) "LATCHED" DEVICE. NOW THE SELECTED INTERRUPT IS "ACKNOWLEDGED" BY THE CHANNEL CONTROLLER. THE SELECTED DEVICE IS GIVEN USE OF THE CHANNEL UNTIL THE INTERRUPT IS "RELEASED" BY THE CPU I/O PROCESS OR THE DMA CONTROLLER, WHICHEVER IS INVOLVED WITH THE ACTIVE INTERRUPT.

THE DATA INTERRUPT AND STATUS INTERRUPT MECHANISMS ARE FULLY INDEPENDENT. ANY OF THE PHASES OF THE INTERRUPT SEQUENCE, OF EITHER TYPE OF INTERRUPT, CAN BE MAINTAINED REGARDLESS OF THE PHASE OF THE OTHER.

THE "RELEASE" OF AN INTERRUPT IS NECESSARY TO REACTIVATE THE SCAN FOR ANOTHER INTERRUPT OF THE SAME TYPE, BY BOTH THE CHANNEL CONTROLLER AND DEVICE CONTROL-

LER. THE "RELEASE" IS CAUSED BY DIFFERENT COMMAND SEQUENCES, BASED ON THE TYPE OF INTERRUPT.

FOLLOWING A:	THE INTERRUPT IS RELEASED:
DATA-IN INTERRUPT	-AFTER EXACTLY TWO RIO SIGNALS HAVE BEEN TRANSMITTED TO THE CHANNEL.
DATA-OUT INTERRUPT	-BY RIO FOLLOWED BY A "DATA AVAILABLE" CHANNEL COMMAND (SEE SECTION 8.3.2).
STATUS INTERRUPT	-BY THE FIRST RIO FOLLOWING A "STATUS REQUEST" CHANNEL COMMAND (SEE SECTION 8.3.2).

THE DEVICE CONTROLLER RECOGNIZES THE APPROPRIATE SEQUENCE, AND PRODUCES THE RELEASE SIGNAL WHICH CLEARS ITS "REQUEST FLAG" AND "LATCH FLAG", AND THE "PENDING FLAG" WITHIN THE CHANNEL CONTROLLER.

8.4.2.4 DIRECT REQUEST MECHANISM

THE DIRECT REQUEST MECHANISM ALLOWS AN ACTIVE CPU PROCESS TO READ THE CONTENT OF THE DEVICE CONTROLLER REGISTERS. A "DIRECT REQUEST" MAY BE MADE AT ANY TIME, EXCEPT WHEN THE CHANNEL IS IN THE MIDDLE OF ONE OF THE FOLLOWING INTERRUPT OPERATIONS:

DATA-IN INTERRUPT:	BETWEEN THE FIRST AND SECOND RIO SIGNAL.
DATA-OUT INTERRUPT:	BETWEEN THE RIO AND "DATA AVAILABLE" COMMAND.
STATUS INTERRUPT:	BETWEEN "STATUS REQUEST" AND THE NEXT RIO.

"DIRECT REQUEST" IS ACCOMPLISHED BY EXECUTING ANY OF THE FOLLOWING DEVICE COMMANDS:

- 1) READ STATUS REGISTER.
- 2) READ WORD COUNT REGISTER.
- 3) READ BUFFER ADDRESS REGISTER.
- 4) READ DEVICE CONTROL WORD A.
- 5) READ DEVICE CONTROL WORD B.

(REFER TO SECTION 8.5.3 FOR COMPLETE INFORMATION ON THE USE OF THESE COMMANDS.)

AFTER SENDING A "DIRECT REQUEST" THE PROGRAM MUST WAIT FOR THE TWO WAY DELAY ALONG THE CHANNEL CABLES (300 NANOSECONDS FOR 75 FOOT STANDARD LENGTHS). THIS MEANS THAT IF AN XIO, EXECUTING A "DIRECT REQUEST", IS ISSUED AT T-PERIOD T(0) THE RIO WHICH GATES THE DATA RETURNED MUST NOT BE SPECIFIED BEFORE T-PERIOD T(5). COMPLETION OF THE "DIRECT REQUEST" SEQUENCE WILL RETURN THE CHANNEL TO EXACTLY THE STATE WHICH EXISTED PRIOR TO "DIRECT REQUEST MODE". AN ADDITIONAL TIME DELAY, EQUIVALENT TO THAT REQUIRED FOR A "DIRECT REQUEST" OPERATION, MUST BE ALLOWED BY THE PROGRAM FOLLOWING THE SEQUENCE BEFORE ALLOWING ANY NEW INTERRUPTS TO TAKE EFFECT. THIS IS REQUIRED IN ORDER TO GUARANTEE THE RETURN OF ANY PREVIOUS INFORMATION TO THE CHANNEL DATA AND ID LINES. REFER TO SECTION 8.5.6 FOR AN EXAMPLE OF "DIRECT REQUEST".

8.4.2.5 DEVICE COMMANDS

EACH DEVICE CONTROLLER HAS A CONTROL UNIT WHICH RECOGNIZES AND EXECUTES "COMMANDS" FROM THE CPU, WHEN ACCOMPANIED BY THE MATCHING "DEVICE SELECT" CODE ASSIGNED TO THE CONTROLLER. DEVICE COMMANDS MAY BE USED TO CLEAR A DEVICE, START AN OPERATION, READ DEVICE CONTROL REGISTERS, ETC.

THE FOLLOWING LIST SPECIFIES STANDARD DEVICE COMMANDS, AS ASSIGNED BY NANODATA TO COVER MOST EXTERNAL DEVICE CONTROL FUNCTIONS:

CODE (OCTAL)	COMMAND	DESCRIPTION
01	CLEAR DEVICE	_CLEAR (ZEROS) ALL REGISTERS, RESETS ALL INTERRUPT FLAGS, AND DISABLES FURTHER INTERRUPTS WITHIN THE DEVICE CONTROLLER. STOPS ANY ACTIVITY WITHIN THE DEVICES ATTACHED. ALSO RESULTS IN THE SAME EFFECT AS THE "DISABLE INTERRUPTS" COMMAND (05, BELOW).
04	ENABLE DEVICE INTERRUPTS	_ENABLES THE GENERATION OF INTERRUPTS, OF ANY OF THE THREE TYPES (DATA-IN, DATA-OUT, STATUS), BY THE DEVICE CONTROLLER. IF AN INTERRUPT IS IN THE "REQUESTED" CONDITION, AT THE TIME OF THIS COMMAND, IT WILL BE GENERATED.
05	DISABLE DEVICE INTERRUPTS	_INHIBITS GENERATION OF ANY DEVICE INTERRUPTS. THIS DOES NOT STOP INTERRUPT "REQUESTS" FROM OCCURRING AT THE DEVICE, BUT DOES BLOCK THE "LATCHING" OF ANY INTERRUPTS.
07	READ STATUS	_A "DIRECT REQUEST" COMMAND (SEE SECTIONS 8.4.2.4 AND 8.5.6) TO READ THE STATUS REGISTER (SR) OF THE SPECIFIED DEVICE CONTROLLER ON THE SELECTED CHANNEL.
11	LOAD WORD COUNT	_LOADS THE VALUE ON THE SELECTED CHANNEL DATA LINES INTO THE WORD COUNT REGISTER (WC) OF THE SPECIFIED DEVICE CONTROLLER.
12	LOAD BUFFER ADDRESS REGISTER	_LOADS THE VALUE ON THE SELECTED CHANNEL DATA LINES INTO THE BUFFER ADDRESS REGISTER (BAR) OF THE SPECIFIED DEVICE CONTROLLER.
21	READ WORD COUNT	_A "DIRECT REQUEST" COMMAND TO READ THE WORD COUNT REGISTER (WC) FROM THE SPECIFIED DEVICE CONTROLLER ON THE SELECTED CHANNEL.

22	READ BUFFER ADDRESS REGISTER	_A "DIRECT REQUEST" COMMAND TO READ THE BUFFER ADDRESS REGISTER (BAR) FROM THE SPECIFIED DEVICE CONTROLLER ON THE SELECTED CHANNEL.
30	LOAD DEVICE CONTROL WORD A	_LOADS THE VALUE ON THE SELECTED CHANNEL DATA LINES INTO DEVICE CONTROL WORD "A" (DCWA) OF THE SPECIFIED DEVICE CONTROLLER.
31	LOAD DEVICE CONTROL WORD B	_LOADS THE VALUE ON THE SELECTED CHANNEL DATA LINES INTO DEVICE CONTROL WORD "B" (DCWB) OF THE SPECIFIED DEVICE CONTROLLER.
34	READ DEVICE CONTROL WORD A	_A "DIRECT REQUEST" COMMAND TO READ DEVICE CONTROL WORD "A" (DCWA) FROM THE SPECIFIED DEVICE CONTROLLER ON THE SELECTED CHANNEL.
35	READ DEVICE CONTROL WORD B	_A "DIRECT REQUEST" COMMAND TO READ DEVICE CONTROL WORD "B" (DCWB) FROM THE SPECIFIED DEVICE CONTROLLER ON THE SELECTED CHANNEL.
41	START OPERATION	_STARTS THE DEVICE OPERATION, GENERATING THE FIRST DATA INTERRUPT CYCLE, IF APPROPRIATE.
46	CLEAR ERROR	_CLEARS ALL ERROR BITS IN THE STATUS REGISTER (SR) AND INTERRUPT STATUS REGISTER (ISR) OF THE SPECIFIED DEVICE CONTROLLER.

8.4.3 DATA ROUTING SUPPORT

8.4.3.1 BUFFER ADDRESS REGISTER (BAR)

THE BUFFER ADDRESS REGISTER SERVES AS A MEMORY POINTER FOR DATA TRANSFER OPERATIONS. IT IS AN 18 BIT WIDE REGISTER WHICH IS LOADED BY THE "LOAD BUFFER ADDRESS REGISTER" DEVICE COMMAND (SEE SECTION 8.4.2.5). THE BUFFER ADDRESS REGISTER HAS BEEN INCORPORATED INTO THE DEVICE CONTROLLER IN ORDER TO REMOVE CPU OVERHEAD DURING DATA TRANSFERS. THE CONTENT OF THE BAR WILL BE PLACED ON THE CHANNEL DATA LINES FOR USE BY THE CPU, OR DMA CONTROLLER, ON THE OCCURRENCE OF ANY DATA INTERRUPTS. THE REGISTER MAY BE INCREMENTED, OR DECREMENTED, AUTOMATICALLY AT THE COMPLETION OF EACH DATA WORD TRANSFER. THE CPU PROGRAM CAN REQUEST THIS ACTION BY ENABLING THE CORRESPONDING DCWB BITS (SEE SECTION 8.4.2.1). BAR MAY ALSO BE READ AT ANY TIME, THROUGH A "DIRECT REQUEST" COMMAND (SEE SECTION 8.4.2.4).

8.4.3.2 WORD COUNT REGISTER (WC)

WORD COUNT IS ANOTHER REGISTER DESIGNED TO MINIMIZE CPU OVERHEAD DURING DATA TRANSFER OPERATIONS. WC WILL NORMALLY CONTAIN THE VALUE REPRESENTING THE LENGTH OF THE MEMORY BLOCK (BUFFER) ADDRESSED BY BAR. IT IS LOADED USING THE "LOAD WORD COUNT" DEVICE COMMAND (SEE SECTION 8.4.2.5). IF ENABLED BY THE APPROPRIATE DCNB BIT (SECTION 8.4.2.1) WC WILL BE DECREMENTED, AUTOMATICALLY, EVERY TIME A WORD TRANSFER HAS BEEN COMPLETED. UPON DECREMENTING TO ZERO THE "WORD COUNT REACHED ZERO" STATUS REGISTER (SR) BIT IS SET. THE EQUIVALENT BIT IN THE INTERRUPT STATUS REGISTER (ISR) WILL BE SET, AND A STATUS INTERRUPT GENERATED, IF THE "WORD COUNT ZERO INTERRUPT" CONDITION IS ENABLED IN DCNR (SEE SECTION 8.4.2.1). FURTHER DATA INTERRUPTS ARE INHIBITED ONCE WC REACHES ZERO. MANY DEVICES MAY ALSO BE ENABLED TO COMPLETE THEIR OPERATING CYCLE AUTOMATICALLY, UPON WC REACHING THIS CONDITION. WC MAY BE READ AT ANY TIME BY A "DIRECT REQUEST" COMMAND (SECTION 8.4.2.4).

THE WORD COUNT REGISTER IS 16 BITS WIDE. IT MAY CONTAIN VALUES BETWEEN 0 AND $(2^{16})-1$. THE NEXT HIGHER BIT (BIT POSITION 16) IS USED TO INDICATE A "DATA CHAINING" MODE OF OPERATION, ON THE CURRENT DATA BLOCK (SEE SECTION 8.4.3.4).

AN I/O DEVICE MAY BE OPERATED WITHOUT USING THE WORD COUNT REGISTER BY SIMPLY LEAVING THE "DECREMENT WORD COUNT" CONTROL (DCNB BIT 6) RESET (0). THE WORD COUNT REGISTER ITSELF MUST BE LOADED WITH A NON-ZERO VALUE, IF DATA INTERRUPTS ARE TO BE ALLOWED TO OCCUR. THE PROGRAMMER MUST BE AWARE THAT THIS BECOMES A HAZARDOUS MODE OF OPERATION, AS THERE IS NOW NO HARDWARE PROTECTION AGAINST A DATA TRANSFER OVERRUNNING ITS MEMORY BUFFER BOUNDARY.

8.4.3.3 DATA TRANSFERS

DATA TRANSFERS ARE MAINTAINED ON A WORD BY WORD BASIS. ALL OF THE UNITS DESCRIBED SO FAR, IN SECTION 8.4, ESTABLISH A SYSTEM WHICH ALLOWS DATA TRANSFERS TO BE PERFORMED WITH LITTLE CPU OVERHEAD. AT THE SAME TIME THAT A DATA INTERRUPT SIGNAL IS BEING SENT TO QM-1, THE BUFFER ADDRESS REGISTER AND TASK ID ARE BEING PRESENTED TO THE I/O CHANNEL. THE PROGRAM RESPONDING TO THE INTERRUPT NEED NOT KNOW WHICH DEVICE ACTUALLY GENERATED THE INTERRUPT.

THE NANOSTORE LOCATION TO WHICH THE DATA INTERRUPT IS ASSOCIATED DETERMINES THE OPERATION TO BE PERFORMED FOR THAT INTERRUPT. SEPARATE NANOSTORE LOCATIONS ARE REQUIRED FOR EACH DATA-IN AND DATA-OUT INTERRUPT POSSIBLE. THE DATA ROUTING PROCEDURE APPLIED MAY SIMPLY IGNORE THE TASK ID AND ROUTE THE DATA WORD TO OR FROM THE MAIN STORE LOCATION IDENTIFIED BY THE CONTENT OF THE BAR PRESENTED TO THE CHANNEL. THE PROCEDURE MAY OPTIONALLY USE THE TASK ID TO DETERMINE TO WHICH OF THE QM-1 MEMORIES THE CURRENT DATA WORD SHOULD BE ROUTED (SUCH AS MAIN STORE, CONTROL STORE, ETC.).

ON DATA-OUT INTERRUPTS, WHERE DATA IS TO BE ROUTED FROM THE CPU TO THE DEVICE, THE PROGRAM IMMEDIATELY ISSUES RIO THEREBY READING BAR INTO THE PORT REGISTER. THIS RIO ALSO SWITCHES THE DEVICE CONTROLLER TO A "WAITING FOR DATA" CONDITION. WHEN THE "DATA AVAILABLE" CHANNEL COMMAND (SECTION 8.3.2) IS SENT, THE DEVICE CONTROLLER ACCEPTS THE DATA THEN PRESENT ON THE DATA LINES. THERE IS NO NEED FOR ANY TIME DELAY BETWEEN THE RIO THAT READS BAR AND THE "DATA AVAILABLE" COMMAND.

IN THE CASE OF DATA-IN INTERRUPTS, THE CHANNEL CONTROLLER EXPECTS TWO CONSECUTIVE RIO SIGNALS TO COMPLETE THE TRANSFER. THE FIRST RIO GATES BAR INTO THE PORT REGISTER AND SIGNALS THE DEVICE CONTROLLER TO PLACE THE DATA WORD ONTO THE CHANNEL DATA LINES. THE CHANNEL NEEDS ABOUT 300 NANOSECONDS (FOR A 75 FOOT CABLE) TO RESPOND TO THE FIRST RIO BEFORE THE DATA IS AVAILABLE FOR THE SECOND RIO. THIS MEANS THAT IF THE FIRST RIO IS SENT AT T-PERIOD T(0) THE SECOND CANNOT BE ISSUED BEFORE T-PERIOD T(5).

8.4.3.4 DATA CHAINING

DATA CHAINING IS A MECHANISM THAT ASSISTS THE I/O PROCESS IN COMBINING SEPARATE BLOCKS OF DATA IN MEMORY INTO ONE LOGICALLY CONSECUTIVE BUFFER SPACE. IN ORDER TO INVOKE THIS MODE OF OPERATION BIT 16 OF THE WORD COUNT REGISTER (SECTION 8.4.3.2) MUST BE SET TO A ONE. THIS IS ACCOMPLISHED AT THE SAME TIME THE ACTUAL WORD COUNT VALUE, DESCRIBING A DATA SEGMENT LENGTH, IS LOADED INTO THE LOWER 16 BIT POSITIONS OF WC USING THE "LOAD WORD COUNT" DEVICE COMMAND (SECTION 8.4.2.5). WITH THE PROPER INTERRUPT ENABLES ALLOWED, WHEN WC REACHES ZERO IT WILL GENERATE A STATUS INTERRUPT AND WILL SET THE "DATA CHAINING" BIT IN THE STATUS REGISTER (SR) AND INTERRUPT STATUS REGISTER (ISR). THE "DATA CHAINING" STATUS BIT REPLACES THE NORMAL SETTING OF THE "WORD COUNT REACHED ZERO" BIT IN THE ISR.

FOLLOWING A "DATA CHAINING" STATUS INTERRUPT THE I/O CONTROL PROCESS MUST FETCH BOTH THE BAR AND WC INFORMATION THAT IDENTIFIES THE NEXT BUFFER SEGMENT. THIS NEW DATA ROUTING INFORMATION IS TRANSFERRED TO THE DEVICE CONTROLLER USING THE "LOAD BUFFER ADDRESS REGISTER" AND "LOAD WORD COUNT" COMMANDS. AS SOON AS THIS IS ACCOMPLISHED A "START OPERATION" COMMAND (SECTION 8.4.2.5) IS ISSUED TO ACTIVATE THE CONTINUATION OF THE DATA TRANSFER, USING THE NEW BAR AND WC VALUES. THE TRANSFER OF THE LAST BUFFER SEGMENT, OF A SET, IS INDICATED BY THE PRESENCE OF A "ZERO" DATA CHAINING MODE BIT IN THE WC REGISTER. SEE SECTION 8.5.7 FOR ADDITIONAL DISCUSSION.

8.4.3.5 DIRECT MEMORY ACCESS (DMA)

DIRECT MEMORY ACCESS (DMA) MODE OF OPERATION PERMITS DIRECT DATA TRANSFER BETWEEN I/O DEVICES AND QM-1 MAIN STORE. NO CPU INTERVENTION IS REQUIRED FOR THE DURATION OF DIRECT MEMORY TRANSFERS. THIS REDUCES DATA HANDLING OVERHEAD FOR THE CPU, AND ALSO PERMITS THE ATTACHMENT TO THE SYSTEM OF VERY HIGH SPEED I/O DEVICES THAT WOULD OTHERWISE BE TOO FAST FOR NORMAL CPU DATA INTERRUPT RESPONSE. DMA ALSO PERMITS HIGHLY EFFICIENT NANOCODING OF NON-INTERRUPTABLE, LONG DURATION, COMPUTATIONAL PROCESSES; SINCE THE CPU IS NOT REQUIRED TO HANDLE FREQUENT DATA INTERRUPTS.

ANY STANDARD DEVICE CONTROLLER MAY BE SWITCHED TO DMA MODE, AS LONG AS THE CHANNEL IS EQUIPPED WITH THE DMA OPTION AND A MULTI-PORT MAIN STORE SYSTEM. TO UTILIZE A DMA DATA PATH, DEVICE OPERATION IS INITIATED IN THE SAME MANNER AS FOR REGULAR CPU DATA TRANSFERS. ALL DEVICE CONTROLLER REGISTERS SHOULD BE LOADED TO INITIALIZE THE TRANSFER AS IF IT WERE TO BE VIA A CPU PORT. THE ONLY DIFFERENCE IS THAT BIT 9 OF DCNB MUST BE SET TO "ONE" (DMA MODE BIT, SEE SECTION 8.4.2.1). WHEN THE DEVICE AND CONTROLLER HAVE BEEN PREPARED FOR THE DATA TRANSFER THE PROGRAM SHOULD ISSUE THE APPROPRIATE "START OPERATION" DEVICE COMMAND (SEE SECTION 8.4.2.5). FROM THIS POINT ON, THE DEVICE CONTROLLER COMMUNICATES DIRECTLY WITH THE MEMORY ACCESS CONTROLLER.

SWITCHING A DEVICE TO DMA MODE DOES NOT DISCONNECT IT FROM THE QM-1 CHANNEL. ALL OF THE FACILITIES TO "MEASURE THE PULSE" OF, OR TO OVERRIDE, THE CURRENT I/O OPERATION REMAIN ACTIVE. FOR EXAMPLE, "DIRECT REQUEST MODE" MAY BE APPLIED AT ANY TIME TO SAMPLE AN ACTIVE REGISTER SUCH AS BAR OR WC. ANY ENABLED STATUS INTERRUPT MAY OCCUR AS USUAL. THE NORMAL TERMINATION OF THE DMA DATA BLOCK TRANSFER WILL USUALLY OCCUR WHEN WC REACHES ZERO. THE CPU MAY BE ASYNCHRONOUSLY NOTIFIED OF THIS EVENT THROUGH THE "WORD COUNT REACHED ZERO" STATUS INTERRUPT, AS INDICATED BY THE CORRESPONDING BIT BEING SET IN THE INTERRUPT STATUS REGISTER. ABNORMAL TERMINATION WILL OCCUR AS A RESULT OF ANY PRESELECTED ERROR INTERRUPT CONDITION. TWO "ERROR" STATUS INDICATORS ARE ASSOCIATED WITH DMA MODE (ALSO REFER TO SECTION 8.4.2.2).

DMA ADDRESSING ERROR - INDICATES AN ATTEMPTED REFERENCE BEYOND THE MEMORY SPACE INSTALLED ON THE COMPUTER SYSTEM. FOR THOSE SYSTEMS EQUIPPED WITH THE VIRTUAL MAIN STORE OPTION THIS BIT INDICATES THAT A REFERENCE TO A NON-RESIDENT USER PAGE WAS ATTEMPTED.

DMA EXCEPTION - INDICATES THAT DMA ENCOUNTERED A MAIN STORE PARITY ERROR, OR SOME OTHER HARDWARE MALFUNCTION.

8.5 I/O SYSTEM OPERATIONS

8.5.1 GENERAL

THIS SECTION DISCUSSES ACTUAL I/O SYSTEM OPERATIONS AND PROCEDURES. SEVERAL ACTUAL NANOCODE EXAMPLES ARE INCLUDED AS AN ILLUSTRATION OF THE VARIOUS CODING TECHNIQUES APPLICABLE IN I/O PROGRAMMING.

8.5.2 INITIALIZATION

CHANNEL INITIALIZATION IS USUALLY PERFORMED ONLY WHEN THE STATE OF EITHER THE CHANNEL OR ITS DEVICES IS UNKNOWN. THIS NORMALLY WILL BE REQUIRED FOLLOWING A SYSTEM POWER UP, OR AFTER A CHANNEL HAS BEEN DEACTIVATED FOR MAINTAINANCE ACTIVITY. INITIALIZATION WILL BE PERFORMED BY OPERATING SYSTEMS DURING SYSTEM INITIAL PROGRAM LOADING AND BY STAND ALONE PROCESSES AT THE START OF EXECUTION.

THIS SECTION DESCRIBES SEVERAL ASPECTS OF INITIALIZING A CHANNEL OR A DEVICE. IT IS MORE SUGGESTIVE THAN IMPERITIVE SINCE THIS OPERATION CAN BE PERFORMED IN SEVERAL WAYS.

THE FIRST STEP OF INITIALIZATION IS TO CLEAR THE CHANNELS AND DEVICES. THREE MEANS OF GENERATING THE CLEAR SIGNAL ARE AVAILABLE:

1. MASTER CLEAR (SEE SECTION 8.2.2) CLEARS ALL CHANNELS AND DEVICES ALONG WITH THE ENTIRE CPU.
2. I/O RESET (SEE SECTION 8.3.2) CLEARS THE CHANNEL AND ALL DEVICES ON THE "KA" SELECTED PORT.
3. CLR COMMAND (SEE SECTION 8.4.2.5) WILL CLEAR THE DEVICE CONTROLLER TO WHICH IT IS SENT.

TWO DEVICE COMMANDS "ENABLE ALL DEVICE INTERRUPTS" AND "DISABLE ALL DEVICE INTERRUPTS" (SEE SECTION 8.4.2.5) ACT AS "CONNECT DEVICE TO THE CHANNEL" AND "DISCONNECT DEVICE FROM THE CHANNEL" RESPECTIVELY. FOLLOWING A CHANNEL OR DEVICE CLEARING OPERATION DEVICE INTERRUPTS ARE LEFT DISABLED. DISABLING THE DEVICE INTERRUPTS REMOVES THE DEVICE FROM THE PRIORITY CHAIN AND DOES NOT ALLOW ANY DEVICE ORIGINATED ACTIVITY. HOWEVER IT IS STILL POSSIBLE TO LOAD ANY REGISTER OF THE DEVICE CONTROLLER AS WELL AS TO READ ANY REGISTER AND TO CHECK STATUS, THROUGH DIRECT REQUEST COMMANDS (SEE SECTION 8.5.6).

IT IS RECOMMENDED THAT EACH DEVICE CONTROL PROGRAM CLEAR THE DEVICE CONTROLLER AND LOAD ITS INITIAL CONTROL REGISTERS BEFORE ISSUING "ENABLE DEVICE INTERRUPTS" AND "START OPERATION" COMMANDS.

8.5.3 COMMAND EXECUTION

COMMANDS (CHANNEL AND DEVICE, SEE SECTIONS 8.3.2 AND 8.4.2.5) ARE GENERATED BY SETTING THE DEVICE NUMBER ON THE "DEVICE SELECT" LINES AND THE COMMAND CODE ON THE "COMMAND" LINES, AND THEN EXECUTING XIO. THE FOLLOWING ARE TWO NANOPROGRAM SEGMENTS TO GENERATE "I/O RESET" (CHANNEL COMMAND) AND TO "LOAD DCWA" (DEVICE COMMAND).

* SEND I/O RESET (70) TO CHANNEL 2.

: "SINCE I/O RESET IS A CHANNEL COMMAND NO DEVICE SELECTION IS NECESSARY"

.... KA = 2 "PORT", KB = 70 "CODE FOR I/O RESET FUNCTION"

X... KB->FIPH "PLACE CODE ON COMMAND LINES"

XIO "TRANSMIT COMMAND"

READ NS "FETCH NEXT NANOWORD"

.X.. GATE NS "PROCEED TO NEXT NANOPROGRAM SEGMENT"

* LOAD DCWA TO DEVICE 17 ON CHANNEL 4.

: "ASSUME DATA TO BE LOADED IS ALREADY IN PORT REGISTER 4"

.... KA = 4 "PORT", KB = 12 "CODE FOR LOAD DCWA", KX = 17 "DEVICE ID"

X... KB->FIPH "PLACE CODE ON COMMAND LINES"

G(G KX) "PLACE DEVICE ID ON DEVICE SELECT LINES (G-BUS)"

XIO "TRANSMIT COMMAND AND DATA"

READ NS "FETCH NEXT NANOWORD"

.X.. GATE NS "PROCEED TO NEXT NANOPROGRAM SEGMENT"

8.5.4 DATA TRANSFERS

DATA TRANSFERS ARE NORMALLY INVOKED BY A DATA INTERRUPT. ON ANY DATA INTERRUPT THE CHANNEL WILL PRESENT THE CONTENTS OF BAR (BUFFER ADDRESS REGISTER) FOR GATING INTO THE PORT. THE ISSUING OF AN RIO ON "DATA IN" INTERRUPTS WILL GATE THE BAR INTO THE PORT REGISTER, SWITCH THE DEVICE TO DATA MODE, AND THEN CAUSE THE DEVICE TO PUT ITS DATA ON THE DATA LINES FOR LATER GATING INTO THE PORT. THE NEXT (SECOND) RIO THEN GATES THE DATA INTO THE PORT REGISTER AND GENERATES THE RELEASE SIGNAL TO THE CHANNEL (SEE TIMING NOTE BELOW). IN THE CASE OF "DATA OUT" INTERRUPTS A "DATA AVAILABLE" CHANNEL COMMAND REPLACES THE SECOND RIO FOR RELEASING THE CHANNEL. UPDATING OF BAR AND NC (WORD COUNT) OCCUR, IF ENABLED, ON THE SECOND RIO FOR THE "IN" CASE AND ON "DATA AVAILABLE" FOR THE "OUT" CASE.

TIMING NOTE: THE FOLLOWING RESTRICTION MUST BE OBSERVED. THE RIO WHICH GATES THE DATA (THE SECOND RIO AFTER A DATA-IN INTERRUPT) MAY NOT BE ISSUED BEFORE THE 5TH T-PERIOD AFTER THE FIRST RIO.

SAMPLE DATA OUT INTERRUPT HANDLER

THE FOLLOWING CODE WILL READ THE BAR OF THE INTERRUPTING DEVICE, USE THAT VALUE TO ADDRESS ONE WORD OF MAIN STORE, READ THAT MS LOCATION AND SEND THE DATA RECEIVED TO THE INTERRUPTING DEVICE, THUS RELEASING THE CHANNEL AND SATISFYING THE DATA-OUT INTERRUPT.

```

*****
*                                     *
*          EQUATES FOR THIS FUNCTION          *
*                                     *
*****

```

```

PORT NUM   = 4           "PORT 4 IDENTITY"
PORT PATH  = PORT NUM + 32 "VALUE 36, FOR ADDRESSING ES(4) AS LS(36)"
DATAVAIL   = 71          "VALUE OF DATA AVAILABLE CHANNEL COMMAND"

```

```

*****
* DATA OUT INTERRUPT HANDLER:                                     *
* ROUTES ONE DATA WORD FROM MAIN STORE TO CHANNEL 4              *
* RETURNS CONTROL TO THE INTERRUPTED PROCESS FOLLOWING TRANSFER    *
*****

```

```

:      "DATA OUT INTERRUPT LEVEL XX, ENTRY POINT"
....  BRANCH(N.+1), KA = PORT_NUM, KB = PORT_PATH, KX = MS_BUSY
S...  GATE NS(NOT X)      "AWAIT MAIN STORE NOT BUSY"
      KB->FMIX, KB->FMD    "SET MAIN STORE PATH TO EXTERNAL STORE"
.X...  RIO               "GATE BAR TO PORT REGISTER, ES(4)"
      READ NS            "PREPARE TO CONTINUE TO NEXT NANOWORD"
..X...  READ MS          "FETCH NECESSARY MAIN STORE WORD"
      GATE NS            "CONTINUE IN NEXT NANOWORD"

:      "COMPLETION OF DATA OUT INTERRUPT LEVEL XX"
....  ALLOW INTS "ON EXIT", KX = MS_DATA, KB = DATAVAIL
S...  GATE NS(NOT X)      "AWAIT MAIN STORE DATA AVAILABLE"
      GATE MS            "GATE MS DATA EACH TIME UNTIL SUCCESSFUL"
.X...  XIO               "SIGNAL CHANNEL CONTROLLER OF INTENTIONS"
      KB->FIPH            "DATA AVAILABLE COMMAND"
      READ NS            "READ NEXT NANOWORD OF INTERRUPTED PROGRAM"
..X...  GATE NS            "CONTINUE INTERRUPTED PROGRAM"

```

8.5.5 STATUS INTERRUPT HANDLING

THE STATUS INTERRUPT MECHANISM OF THE DEVICE CONTROLLER IS DESCRIBED IN SECTION 8.4.2.2. WHEN THE CPU SENSES A STATUS INTERRUPT IT MAY ISSUE A "STATUS REQUEST" CHANNEL COMMAND. THE INTERRUPT STATUS WILL BE AVAILABLE FOR GATING INTO THE PORT AFTER THE 5TH T-PERIOD FOLLOWING THE STATUS REQUEST. A STATUS REQUEST MAY NOT BE ISSUED IN THE MIDDLE OF DATA INTERRUPT HANDLING, I.E. AFTER THE RIO THAT READ BAR INTO THE PORT.

THE FOLLOWING NANOWORD WILL READ THE INTERRUPT STATUS OF THE INTERRUPTING DEVICE ON CHANNEL 4 INTO LS(G.STAT). WHERE G.STAT REPRESENTS A G-REGISTER THAT POINTS TO THE LOCAL STORE REGISTER TO RECEIVE THE STATUS. G.DEV IS USED TO REPRESENT ANOTHER G-REGISTER WHICH WILL ITSELF RECEIVE THE DEVICE ID RETURNED BY THE INTERRUPTING DEVICE.

```
* THE STATUS INTERRUPT HANDLER WILL TRANSFER CONTROL TO THE SYSTEM SUPERVISOR
* TO ANALYSE THE STATUS RETURN, AND TO CHANGE THE OPERATING STATE IF NEEDED.
: "STATUS INTERRUPT LEVEL YY, ENTRY POINT"
.... BRANCH(SYSTEM INT) "SUPERVISOR TRANSFER ADDRESS"
      KA = 4 "PORT", KB = 77 "STATUS REQUEST COMMAND"
S... XIO, KB->FIPH "PLACE CHANNEL IN STATUS MODE"
.S... G(G.STAT) "IDENTIFY G REGISTER"
      G->FEOD, KA->FEOA "SET UP EXTERNAL STORE TO LOCAL STORE PATH"
..X. IO ID->G.DEV "SAVE THE DEVICE ID, WHICH IS NOW AT THE PORT"
...S RIO, GATE ES "READ THE PORT AND PASS THE STATUS TO LS(G.STAT)"
      READ NS, GATE NS "TRANSFER CONTROL TO THE SYSTEM SUPERVISOR"
```

8.5.6 DIRECT REQUEST COMMANDS

THE DIRECT REQUEST COMMANDS ALLOW A QM-1 PROGRAM TO EXAMINE ALL OF THE CONTROL REGISTERS OF A DEVICE WITHOUT AFFECTING THE CURRENT STATE OF THE CHANNEL OR DEVICE. AFTER SENDING A DIRECT REQUEST COMMAND THE CPU PROGRAM MUST WAIT FOR A MINIMUM OF 5 T-PERIODS BEFORE GATING ANY EXPECTED DATA INTO THE PORT. THE PROGRAM MUST THEN SEND AN RIO TO READ THE DESIRED DATA INTO THE QM-1 PORT REGISTER. THE RIO ALSO RETURNS THE CHANNEL TO WHATEVER STATE IT WAS IN PRIOR TO THE DIRECT REQUEST. FOR THOSE DIRECT REQUEST COMMANDS THAT OUTPUT DATA NO TIME DELAY IS REQUIRED. IN THIS CASE THE DIRECT REQUEST TRANSMITS THE COMMAND AND ANY DATA IMMEDIATELY. AN RIO IS NOT ALLOWED, AND IS NOT EXPECTED BY THE CHANNEL.

AN ADDITIONAL TIMING NOTE: DATA INTERRUPTS OCCUR AT THE MOMENT WHEN THE BAR DATA IS IMMEDIATELY AVAILABLE AT THE PORT DATA LINES, AND THE "TASK ID" IS AVAILABLE ON THE DEVICE ID LINES. THIS INFORMATION IS EXPECTED BY THE DATA INTERRUPT HANDLERS. THEREFORE, WHEN A DIRECT REQUEST HAS TEMPORARILY CHANGED THE INFORMATION CONTENT ON THOSE LINES (IE. AN "INPUT WORD COUNT REGISTER" COMMAND) THE PROCESS IN CONTROL MUST GUARANTEE AT LEAST ANOTHER 5 T-PERIODS (300 NANoseconds), AFTER RELEASE OF DIRECT REQUEST MODE, TO ALLOW RETURN OF THE ORIGINAL DATA-MODE INFORMATION.

A SIMPLE NANOPROGRAM TO READ THE DCWA REGISTER MIGHT BE AS FOLLOWS.

```
RDCW:  "DCWA OF DEVICE B ON PORT 4 IS READ INTO THE PORT REGISTER"
....  KA = 4 "PORT", KB = 34 "COMMAND TO READ DCWA"
S...  XIO, G(G B), KB->FIPH "SEND DIRECT REQUEST TO DEVICE B"
.S...
...X.  READ NS                      "WAIT AT LEAST 5 T-PERIODS"
...X.  RIO, GATE NS                 "READ INTO PORT REGISTER AND EXIT"
```


8.5.7 DATA CHAINING

DATA CHAINING IS A FACILITY WHICH ENABLES THE PROGRAM TO CONNECT SEPARATE BLOCKS OF DATA IN MAIN MEMORY AS ONE "CONSECUTIVE" BUFFER. THE DEVICE CONTROLLER CONTAINS BASIC LOGIC TO SUPPORT THIS FACILITY. SEVERAL OPTIONAL METHODS EXIST, EACH PROVIDING VARYING DEGREES OF EFFICIENCY FOR EFFECTING DATA CHAINING, THOUGH ONLY THE MOST ELEMENTARY QM-1 METHOD IS DESCRIBED.

TO PERMIT DATA CHAINING THE SUPPORTING PROGRAM MUST PROVIDE A LIST OF BUFFER ADDRESSES AND DATA LENGTHS IN ADVANCE OF STARTING THE CHAINED OPERATION. INITIALLY THE FIRST BUFFER SEGMENT STARTING ADDRESS IS LOADED INTO THE BAR, AND THE WORD COUNT PLUS DATA CHAINING INDICATOR (BIT 16 OF THE WORD COUNT VALUE) IS LOADED INTO WC. THE DEVICE CONTROLLER WILL SENSE WC GOING ZERO AND GENERATE A STATUS INTERRUPT. IN THIS CASE THE INTERRUPT STATUS REGISTER WILL BE FOUND TO HAVE BIT 5 (DATA CHAINING IN EFFECT) SET IN PLACE OF THE USUAL BIT 4 (WORD COUNT GOING ZERO). THE STATUS INTERRUPT HANDLER MUST THEN FETCH THE NEXT SET OF WORDS IDENTIFYING THE FOLLOWING BUFFER SEGMENT, FROM THE BUFFER LIST, AND TRANSFER THEM TO THE BAR AND WC. THIS IS FOLLOWED BY A "START" COMMAND WHICH CONTINUES THE ORIGINAL DATA TRANSFER OPERATION WITH A DATA INTERRUPT.

DATA CHAINING PROCEDURES FOR DEVICES WITH HIGH TRANSFER RATES REQUIRE MORE AUTOMATED METHODS OF CHAINING. THESE METHODS ARE UNIQUE TO EACH HIGH SPEED DEVICE, AND ARE DISCUSSED INDEPENDENTLY WITHIN THEIR RESPECTIVE DEVICE CONTROLLER MANUALS.

8.6 EXAMPLE OF AN ACTUAL DEVICE CONTROLLER

8.6.1 DEVICE SPECIFICATIONS

THIS SECTION DESCRIBES THE DEVICE CONTROLLER AND OPERATION OF THE NANODATA LP135 LINE PRINTER. THE LP135 IS A LOW SPEED DEVICE CAPABLE OF OUTPUTTING HARD COPY COMPRISED OF PRINTED LINES, UP TO 132 CHARACTERS WIDE, AT A RATE OF 135 LINES PER MINUTE. THERE ARE 64 CHARACTER CODES AVAILABLE, CORRESPONDING TO USASCII CODES 40 TO 137 (OCTAL). THE LP135 ALSO HAS A FOUR CHANNEL VERTICAL FORMAT TAPE.

THE PRINTER HAS A FULL 132 CHARACTER LINE BUFFER. TO PRINT A LINE THE CHARACTERS ARE FIRST LOADED INTO THE LINE BUFFER. THEN, A PRINT SIGNAL IS ISSUED CAUSING THE LINE TO BE PRINTED AND THE FORMS ADVANCED TO THE NEXT LINE. WHILE THE FORMS ARE ADVANCING THE LINE BUFFER MAY BE RE-LOADED, IN ORDER TO BE READY FOR THE NEXT PRINT CYCLE WHEN THE FORMS ARE IN POSITION. IF LESS THAN 132 CHARACTERS HAVE BEEN SENT TO THE BUFFER, WHEN THE PRINT SIGNAL IS ISSUED, THE BUFFER IS AUTOMATICALLY FILLED OUT WITH BLANKS BEFORE THE ACTUAL PRINT OPERATION TAKES PLACE.

THE LP135 ALSO HAS PROVISION FOR FORM FEED CONTROLS, AND FOUR FORMAT CHANNEL INDICATORS. THE TERM "SKEW" IS USED, IN THIS DISCUSSION, TO INDICATE FORMS MOTION THROUGH THE PRINTER. "SKEW MODE" REPRESENTS THE CONTINUOUS FORMS MOTION THROUGH THE PRINTER, BEYOND ONE INTEGRAL LINE. THE NEXT SECTIONS DESCRIBE THE CONTROL OF THE LP135 IN CONJUNCTION WITH THE STANDARD DEVICE CONTROLLER ON THE QM-1.

8.6.2 CONTROLLER SPECIFICATIONS

THE FOLLOWING LISTS ARE DESCRIPTIONS OF THE PRINTER STATUS REGISTERS, DEVICE CONTROL WORDS, AND THE PRINTER DEVICE COMMANDS. THE LISTS SPECIFY STANDARD AS WELL AS DEVICE DEPENDENT FUNCTIONS.

8.6.2.1 DEVICE STATUS

STATUS REGISTER USAGE (SR AND ISR)

REFER TO SECTION 8.4.2.2 FOR A GENERAL DISCUSSION OF THE STATUS REGISTERS.

BIT POSITION & IDENT	DESCRIPTION OF STATUS	APPEARS IN ISR
0 DEVICE READY	_A "1" INDICATES THAT POWER IS APPLIED TO THE PRINT UNIT, PAPER IS PROPERLY INSTALLED, AND THE READY BUTTON HAS BEEN DEPRESSED.	YES
1 DEVICE NOT BUSY	_A "1" INDICATES THE DEVICE IS MOTIONLESS, AND IS NOT IN OPERATION ON ANY COMPUTER ISSUED COMMANDS.	YES
2 ERROR CONDITION	_THE LOGICAL "OR" OF STATUS REGISTER BITS 12, 13, 16, AND 17.	YES
3 BUFFER READY	_THE PRINTER'S LINE BUFFER IS READY FOR CHARACTER LOADING (SEE SECTION 8.4.2.2).	YES
4 WORD COUNT ZERO	_(DESCRIBED IN SECTION 8.4.2.2).	YES
5 DATA CHAINING	_(DESCRIBED IN SECTION 8.4.2.2).	YES
6 LINE BUFFER FULL	_THE 132 CHARACTER LINE BUFFER IS FULL. A "REAL TIME" INDICATION OF LINE BUFFER STATE. THIS CONDITION IS PREREQUISITE TO AN ACTUAL PHYSICAL PRINT CYCLE, BUT IS AUTOMATICALLY PRODUCED FOLLOWING A "PRINT" COMMAND TO THE DEVICE CONTROLLER.	YES
7 LINE COUNT	_A STATUS INTERRUPT PRODUCED, ONLY IF ENABLED, EVERY TIME THE PRINTER PASSES THE START OF A LINE POSITION DURING FORMS MOTION. IF THE PRINTER IS IN "SLEW MODE" THE PROGRAM HAS ONE MILLISECOND TO STOP THE "SLEW" IN ORDER TO STOP ON THE CURRENT LINE. THIS BIT IS RESET BY THE "CLEAR ERROR" COMMAND (SEE SECTION 8.4.2.5).	YES
8 FORMAT CHANNEL 1	_A STATUS INTERRUPT PRODUCED, ONLY IF ENABLED,	YES

EVERY TIME FORMAT CHANNEL 1 IS ENCOUNTERED, ON THE FORMS CONTROL TAPE, DURING CARRIAGE MOTION. IF IN "SLEW MODE", THE PROGRAM HAS ONE MILLI-SECOND TO STOP THE "SLEW" IN ORDER TO STOP ON THE DESIRED LINE. THE BIT IS CLEARED WHEN THE NEXT LINE IS ENCOUNTERED.

9	FORMAT CHANNEL 2	_SAME AS ABOVE, FOR CHANNEL 2.	YES
10	FORMAT CHANNEL 3	_SAME AS ABOVE, FOR CHANNEL 3.	YES
11	FORMAT CHANNEL 4	_SAME AS ABOVE, FOR CHANNEL 4.	YES
12	DMA ADDRESSING ERROR	_STANDARD (SEE SECTION 8.4.2.2).	NO
13	DMA EXCEPTION	_STANDARD (SEE SECTION 8.4.2.2).	NO
16	LOST DATA	_A STATUS INTERRUPT PRODUCED, ONLY IF ENABLED, WHEN AN ATTEMPT HAS BEEN MADE TO LOAD THE LINE BUFFER WHEN THE BUFFER IS NOT YET READY FOR A NEW CHARACTER. THIS DEVICE DOES NOT DEMAND THAT CHARACTERS BE TRANSFERRED TO ITS LINE BUFFER WITHIN A LIMITED TIME FRAME. THEREFORE, THIS ERROR USUALLY WILL INDICATE A HARDWARE MALFUNCTION.	NO
17	COMMAND REJECT	_A STATUS INTERRUPT IS GENERATED IF "COMMAND REJECT" IS ENABLED, WHEN ANY OF THE FOLLOWING CONDITIONS OCCUR: 1) ANY FORM FEED COMMAND ISSUED WHEN THE UNIT IS NOT AVAILABLE (SEE STATUS BIT 3, ABOVE) OR IS NOT READY (STATUS BIT 0). 2) AN ATTEMPT TO CLEAR THE LINE BUFFER WHILE A PRINT CYCLE IS IN PROGRESS. 3) AN ATTEMPT TO PRINT A LINE WITHOUT ONE OF THE FOLLOWING STATUS CONDITIONS BEING TRUE: A) BUFFER READY (UNIT AVAILABLE). OR B) BUFFER FULL.	YES

8.6.2.2 DEVICE COMMANDS

PRINTER DEVICE COMMANDS

REFER TO SECTION 8.4.2.5 FOR A COMPLETE DISCUSSION ON DEVICE COMMANDS. ALL COMMAND CODE VALUES ARE SHOWN IN OCTAL.

CODE	COMMAND	DESCRIPTION
01	CLEAR DEVICE	_STANDARD (SEE SECTION 8.4.2.5).
04	ENABLE DEVICE	_STANDARD.
	INTERRUPTS	
05	DISABLE DEVICE	_STANDARD.
	INTERRUPTS	
07	READ STATUS	_STANDARD.
11	LOAD WORD COUNT	_STANDARD.
12	LOAD BUFFER	_STANDARD.
	ADDRESS REGISTER	
21	READ WORD COUNT	_STANDARD.
22	READ BUFFER	_STANDARD.
	ADDRESS REGISTER	
30	LOAD DEVICE	_STANDARD.
	CONTROL WORD A	
31	LOAD DEVICE	_STANDARD.
	CONTROL WORD B	
34	READ DEVICE	_STANDARD.
	CONTROL WORD A	
35	READ DEVICE	_STANDARD.
	CONTROL WORD B	
40	START OPERATION	_THIS COMMAND COMBINES THE FUNCTIONS OF COMMANDS 41 AND 42, DESCRIBED BELOW. IT CAUSES AUTOMATIC PHYSICAL LINE PRINTING TO OCCUR, AS WOULD BE EFFECTED BY COMMAND 42, IMMEDIATELY FOLLOWING THE WORD COUNT REGISTER REACHING ZERO. IN GENERAL, THIS COMMAND INITIATES ONE FULL LINE PRINTER OPERATING CYCLE; FROM DATA TRANSFER TO COMPLETION OF PHYSICAL LINE PRINTING, AND AUTOMATIC ADVANCE TO THE NEXT. REFER TO COMMANDS 41 AND 42 FOR ADDITIONAL DETAIL.
41	START OPERATION	_ASSUMING THE PRINTER IS READY, AND THE DEVICE CONTROLLER REGISTERS ARE PROPERLY INITIALIZED, WITH NC CONTAINING A NON-ZERO VALUE, A DATA BLOCK TRANSFER

WILL BE STARTED. "START OPERATION" TRIGGERS THE FIRST DATA-OUT INTERRUPT REQUEST TO THE CPU OR DMA (DEPENDING ON DMA MODE SETTING), FOR THE CURRENT DATA BLOCK.

- 42 PRINT _THE PRINTER WILL PRINT THE CURRENT CONTENTS OF THE LINE BUFFER AND THEN AUTOMATICALLY ADVANCE TO THE NEXT LINE. IF THE BUFFER IS NOT FULL IT WILL BE AUTOMATICALLY BLANK FILLED. THIS COMMAND MAY NOT BE ISSUED UNLESS EITHER THE "BUFFER READY" OR "BUFFER FULL" CONDITION IS TRUE (SEE STATUS ABOVE).
- 45 CLEAR BUFFER _BLANK FILLS THE ENTIRE PRINTER LINE BUFFER, AND SETS THE BUFFER TO AN EMPTY CONDITION. MAY BE ISSUED ONLY IF THE "BUFFER READY" STATUS BIT IS TRUE.
- 46 CLEAR ERROR _STANDARD (SEE SECTION 8.4.2.5).

NOTE: ALL OF THE FORMS MOTION (CONTROL) COMMANDS DESCRIBED BELOW REQUIRE THE "BUFFER READY" STATUS CONDITION.

- 50 SPACE
 (LINE FEED) _CAUSES THE PRINTER TO ADVANCE THE FORMS IMMEDIATELY TO THE NEXT LINE.
- 51 SKIP TO FORMAT
 CHANNEL 1 _START FORMS "SLEW". STOP WHEN FORMAT CHANNEL 1 IS ENCOUNTERED.
- 52 SKIP TO FORMAT
 CHANNEL 2 _SAME AS 51, FOR CHANNEL 2.
- 53 SKIP TO FORMAT
 CHANNEL 3 _SAME AS 51, FOR CHANNEL 3.
- 54 SKIP TO FORMAT
 CHANNEL 4 _SAME AS 51, FOR CHANNEL 4.
- 64 START SLEW _STARTS "SLEW MODE" CONTINUOUS FORMS MOTION. MOTION CONTINUES UNTIL A "STOP SLEW" (65), "CLEAR DEVICE" (01), OR "I/O RESET" (70) COMMAND IS ISSUED.
- 65 STOP SLEW _STOPS FORMS MOTION, PREVIOUSLY ISSUED BY ANY OF THE FORMS CONTROL COMMANDS (ABOVE).

8.6.2.3 DEVICE CONTROL WORDS

PRINTER DEVICE CONTROL WORDS (DCWA AND DCWB)

REFER TO SECTION 8.4.2.1 FOR A DISCUSSION OF STANDARD DCW FUNCTIONS.

BIT IDENTITY	DESCRIPTION

DATA MAPPING CONTROLS	
DCWA 02 TRANSLATE	_IF SET, THE PRINTER CONTROLLER WILL TRANSLATE LOWER CASE ASCII CHARACTERS TO UPPER CASE.
DCWA 03 PACK/UNPACK	_IF SET, THE PRINTER CONTROLLER WILL UNPACK TWO 8-BIT BYTES FROM EACH 18 BIT WORD TRANSFERRED. THE LEFTMOST CHARACTER (BIT POSITIONS 8 TO 15) IS ACCESSED FIRST. THE CHARACTERS WITHIN THE 18 BIT WORD ARE RIGHT JUSTIFIED. THE HIGH ORDER TWO BITS ARE IGNORED. IF THE WORD COUNT REGISTER IS BEING EMPLOYED DURING "PACKED" MODE OPERATION IT WILL BE DECREMENTED BY TWO FOR EACH QM-1 DATA WORD TRANSFERRED, THEREBY BECOMING A CHARACTER COUNT.
STANDARD STATUS INTERRUPT ENABLE MASKS	
DCWA 04 COMMAND REJECT	_STANDARD (SEE SECTION 8.4.2.1).
DCWA 05 ANY ERROR CONDITION	_STANDARD.
DCWA 06 DEVICE READY	_ENABLES A STATUS INTERRUPT WHENEVER THE PRINTER READY OR RESET BUTTONS CAUSE A CHANGE IN THE PRINTER'S READY CONDITION.
DCWA 07 NOT BUSY	_ENABLES A STATUS INTERRUPT WHENEVER THE PRINTER COMPLETES A FULL OPERATING CYCLE. INDICATES AN END OF "SLEW MODE" OPERATION AS WELL AS AN ACTUAL PRINT CYCLE.
DCWA 08 UNIT AVAILABLE (BUFFER READY)	_ENABLES A STATUS INTERRUPT WHEN THE LINE BUFFER BECOMES READY FOR CHARACTER LOADING.
DCWA 09 WORD COUNT ZERO	_STANDARD.

PRINTER DEPENDENT STATUS INTERRUPT ENABLE MASKS

DCNB 10	FORMAT CHANNEL 1	_ENABLES A STATUS INTERRUPT EACH TIME FORMAT CHANNEL 1 IS ENCOUNTERED.
DCNB 11	FORMAT CHANNEL 2	_SAME AS DCNB 10, FOR CHANNEL 2.
DCNB 12	FORMAT CHANNEL 3	_SAME AS DCNB 10, FOR CHANNEL 3.
DCNB 13	FORMAT CHANNEL 4	_SAME AS DCNB 10, FOR CHANNEL 4.
DCNB 14	LINE COUNT	_ENABLES A STATUS INTERRUPT EVERY TIME A NEW LINE POSITION IS PASSED, DURING FORMS MOTION.
DCNB 15	BUFFER FULL	_ENABLES A STATUS INTERRUPT WHEN THE PRINTER LINE BUFFER BECOMES FULL.

ALL DCNB FUNCTION FOR THIS DEVICE ARE STANDARD. REFER TO SECTION 8.4.2.1 FOR THE DISCUSSION OF THE DCNB CONTROLS.

8.6.3.1 DATA TRANSFER OPERATION

INITIALIZATION:

THE GENERAL INITIALIZATION PROCEDURE IS DESCRIBED EXTENSIVELY IN SECTION 8.5.2. THE FOLLOWING IS SPECIFIC TO THIS PRINT OPERATION. ASSUME THAT THE PRINTER IS KNOWN TO BE READY, AND OPERATIONAL. THE DCW REGISTERS ARE LOADED.

THE FOLLOWING DCWA STATUS INTERRUPT ENABLE MASKS ARE SET:

DCWA 05 ANY ERROR CONDITION
DCWA 06 DEVICE READY
DCWA 07 NOT BUSY
DCWA 08 BUFFER READY
DCWA 09 WORD COUNT ZERO

THE FOLLOWING DCWB DATA ROUTING SUPPORT CONTROLS ARE SET:

DCWB 06 DECREMENT WORD COUNT
DCWB 07 INCREMENT BUFFER ADDRESS REGISTER

ALL OTHER DCW BITS ARE LEFT RESET (ZERO). BAR IS THEN LOADED WITH THE FIRST WORD ADDRESS OF THE DATA BLOCK IN MAIN STORE. WC IS LOADED WITH THE NUMBER OF WORDS (CHARACTERS) FOR THE LINE TO BE PRINTED. THE PRINTER IS NOW READY FOR THE START OF ITS FIRST OPERATION, WHICH IS DESCRIBED UNDER DATA TRANSFER BELOW.

STATUS:

THE DEVICE STATUS CAN BE READ AT ANY TIME VIA THE "DIRECT REQUEST" MECHANISM (SEE SECTION 8.4.2.4). A STATUS INTERRUPT MAY OCCUR ON ANY OF THE CONDITIONS ENABLED BY DCWA AND DCWB. STATUS INFORMATION ACCESSING OPERATIONS ARE STANDARD, AND ARE EXTENSIVELY DESCRIBED IN SECTION 8.4.2.2.

DATA TRANSFERS:

AFTER THE PRINTER CONTROLLER HAS BEEN INITIALIZED, AS DESCRIBED ABOVE, THE STATUS SHOULD BE READ TO MAKE SURE THAT NO ABNORMAL CONDITION EXISTS AND THAT

THE PRINTER IS READY FOR THE TRANSFER: DEVICE READY, NOT BUSY, BUFFER READY. THE PROGRAM MUST ISSUE AN "ENABLE DEVICE INTERRUPTS" (04) COMMAND WHICH LINKS THE DEVICE INTO THE CHANNEL PRIORITY CHAIN, AND ALLOWS INTERRUPTS TO BE TRANSMITTED TO THE CPU.

THE "START OPERATION" (41) COMMAND WILL GENERATE THE FIRST DATA-OUT INTERRUPT. DATA TRANSFERS ARE FURTHER DESCRIBED IN SECTIONS 8.4.3.3 AND 8.5.4. AFTER EACH WORD (CHARACTER) HAS BEEN RECEIVED BY THE PRINTER CONTROLLER, AND TRANSFERRED TO THE LINE BUFFER, THE CONTROLLER WILL UPDATE THE VALUES IN BAR AND WC AND THEN WILL GENERATE THE NEXT DATA-OUT INTERRUPT. WHEN WC REACHES ZERO A STATUS INTERRUPT IS GENERATED AND THE CORRESPONDING SR AND ISR BITS ARE SET.

BEFORE THE PHYSICAL PRINT CYCLE IS INITIATED, THE PROGRAM SHOULD CHECK WHETHER THE PRINTER IS READY FOR THE "PRINT" (42) COMMAND (SEE THE COMMAND LIST IN SECTION 8.6.2.2, ABOVE). THE "PRINT" COMMAND IS THEN ISSUED. THE END OF THE PRINT CYCLE WILL RESULT IN BOTH THE "BUFFER READY" AND "NOT BUSY" STATUS INDICATIONS OCCURRING TOGETHER (UNIQUE TO THE LP135 PRINTER). FOLLOWING THIS STATUS INTERRUPT THE PRINTER CONTROLLER MAY BE PREPARED FOR ITS NEXT OPERATION.

8.6.3.2 DEVICE CONTROL FUNCTIONS

ALL FORMS CONTROL COMMANDS MAY BE ISSUED ONLY WHEN THE PRINTER "BUFFER READY" STATUS IS SET ("1"). DATA TRANSFERS TO THE LINE BUFFER MAY BE PERFORMED EVEN WHILE A FORM FEED, OR OTHER "SLEW MODE", OPERATION IS IN PROGRESS, IN PREPARATION FOR AN IMMEDIATELY FOLLOWING PRINT CYCLE.

A "SPACE" (50) COMMAND WILL ADVANCE THE PRINTER EXACTLY ONE LINE POSITION.

A "SKIP TO FORMAT CHANNEL 1" (51) COMMAND WILL ADVANCE THE FORMS UNTIL FORMAT CHANNEL 1 IS ENCOUNTERED. SIMILARLY, COMMANDS FOR FORMAT CHANNELS 2, 3, AND 4 WILL HAVE THE SAME AFFECT.

ANOTHER FORM FEED OPERATION IS "SLEW MODE". IT IS INITIATED BY A "START SLEW" (64) COMMAND. THE PROGRAM MAY CONTROL THE AMOUNT OF FORMS ADVANCE BY REFERRING EITHER THE FORMAT CHANNELS OR THE LINE COUNT INDICATORS (SEE SECTION 8.6.2.1, ABOVE). THE "SLEW" IS LIMITED TO FOUR CONVENTIONAL PAGES NORTH OF FORMS AT A TIME, TO AVOID ACCIDENTAL WASTE OF PAPER. AN AUTOMATIC "STOP SLEW" IS ISSUED BY THE CONTROLLER ON A RUNNWAY FORMS CONDITION. TO STOP THE "SLEW" THE PROGRAM MUST ISSUE A "STOP SLEW" (65) COMMAND, OR SOME FORM OF DEVICE CLEARING COMMAND. TO STOP MOTION AT THE LINE POSITION MOST RECENTLY INDICATED THE "STOP SLEW" COMMAND MUST BE ISSUED WITHIN ONE MILLISECOND OF NOTIFICATION.


9 INSTALLATION PLANNING ***PRELIMINARY INFORMATION FOR GUIDELINE USE ONLY***

9.1 GENERAL

THIS SECTION HAS BEEN PREPARED TO SERVE AS GUIDE FOR SITE PLANNING AND SITE PREPARATION FOR THE QM-1 COMPUTING SYSTEM. IT CONTAINS THE GENERAL INFORMATION FOR DETERMINING FLOOR SPACE, AIR CONDITIONING AND POWER REQUIREMENTS.

9.2 SPACE REQUIREMENTS

IN THE STANDARD CONFIGURATION, THE QM-1 CPU CONSISTS OF 3 BAYS. THE 3 BAYS ARE NORMALLY ASSEMBLED INTO A "Y" CONFIGURATION AS SHOWN BELOW:



BAY 1 IS 23.50" WIDE, 27" DEEP AND 61.75" HIGH
 BAY 2 AND BAY 3 ARE BOTH:
 27.75" WIDE, 27" DEEP AND 61.75" HIGH

BAY 2 HOUSES THE CENTRAL PROCESSOR, CONTROL STORE, AND NANOSTORE.
 BAY 3 HOUSES MAIN STORE AND THE POWER SUPPLIES.

BAY 1 IS OPTIONAL, AND NORMALLY HOUSES DISK DRIVES, TAPE DRIVES, AND CONTROLLERS.

ACCESS TO THE INTERNAL COMPONENTS OF THE SYSTEM IS THROUGH SWING OUT DOORS ON EACH SIDE OF BAY 2 AND BAY 3, AND THROUGH THE FRONT OF BAY 1.

BAY 1 MAY CONTAIN ONE DISK DRIVE AND CONTROLLER (NANODATA DD50), AND TWO 12.5 IPS TAPE DRIVES AND CONTROLLER (NANODATA TT12.5). ALTERNATIVELY, BAY 1 MAY CONTAIN TWO DISK DRIVES AND CONTROLLER. IN EITHER CASE, A PULLOUT RACK EXTENDS 23" FOR SERVICE. THE REMAINING SERVICE REQUIREMENTS ARE HANDLED BY REMOVING THE SIDE PANELS.

THE COMPLETE SYSTEM, WITH SWING OUT DOORS FULLY OPENED AND THE RACK FULLY EXTENDED, MEASURES 124" WIDE AND 112" DEEP.

9.3 ENVIRONMENTAL REQUIREMENTS

TEMPERATURE

THE AMBIENT TEMPERATURE OF THE INSTALLATION SITE SHOULD BE MAINTAINED BETWEEN 60 DEGREES F AND 80 DEGREES F. THE RECOMMENDED TEMPERATURE IS 70 DEGREES F. OPERATION WITHIN THE LIMITS GIVEN WILL HAVE NO ADVERSE EFFECT ON SYSTEM PERFORMANCE.

HUMIDITY

A RELATIVE HUMIDITY OF 40 - 60 PERCENT AT THE INSTALLATION SITE IS RECOMMENDED. HUMIDITY CONDITIONS MUST NOT ALLOW CONDENSATION TO OCCUR ON ANY SURFACE OR COMPONENT WITHIN THE SYSTEM. EXCESSIVELY LOW HUMIDITY (IE LESS THAN 25% RH) MAY CAUSE PROBLEMS WITH PRINTER PAPER AND CARDS, DUE TO STATIC CHARGES.

AIR CONDITIONING REQUIREMENTS

THE REQUIREMENTS FOR AIR CONDITIONING WILL VARY GREATLY WITH SYSTEM CONFIGURATION, USE, AND LOCAL CONDITIONS. A TYPICAL INSTALLATION WILL PRODUCE A THERMAL LOAD OF 20,000 TO 25,000 BTU/HR.

9.4 POWER REQUIREMENTS

NOMINAL POWER REQUIREMENTS OF THE QM-1 COMPUTING SYSTEM ARE:

120/208 VOLTS, 3 PHASE,

60 HZ @ 7 KVA.

THE SYSTEM IS ALSO AVAILABLE (ON SPECIAL ORDER) TO OPERATE ON 50 HZ.

UNLESS OTHER POWER CABLING IS REQUESTED, NANODATA WILL SUPPLY THE SYSTEM WITH A POWER CABLE TERMINATED IN A MALE PLUG AS FOLLOWS:

3 POLE, 4 WIRE GROUNDING, 125/250 VOLTS.

HUBBELL PLUG 14-30P, MODEL NUMBER 9431 OR 9432

THE POWER CABLE MAY ENTER BAY 3 AT EITHER THE BOTTOM OR TOP.

THE CUSTOMER'S SITE SHOULD PROVIDE THE FOLLOWING POWER:

ONE LINE, 3 PHASE, 120/208 VOLTS, 4 WIRE @ 40 AMPS, WITH

RECEPTICLE - HUBBELL 14-30R, MODEL NUMBER 9430 OR EQUIVALENT.

CIRCUIT BREAKER PANEL WITH

1 CIRCUIT--3 PHASE/30 AMPS

9 SINGLE PHASE CIRCUITS (10 AMP CIRCUIT BREAKERS - 3 PER PHASE

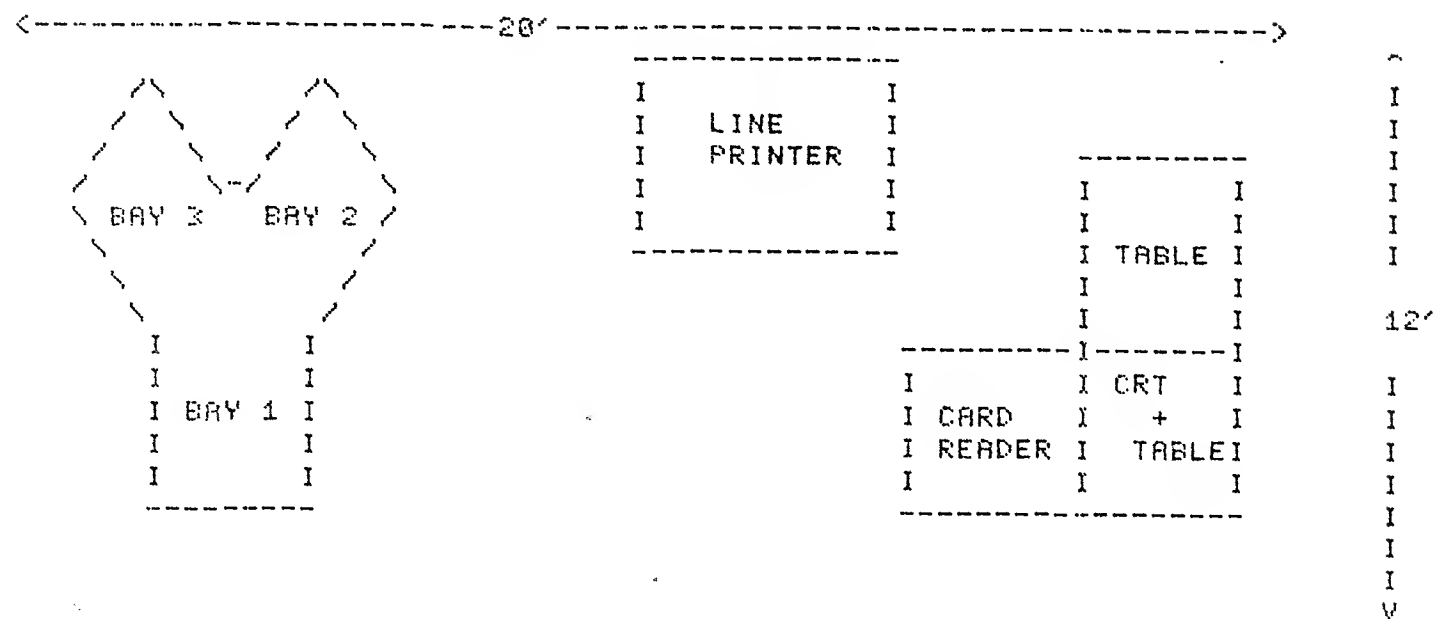
9.5 PERIPHERAL REQUIREMENTS

THE SPACE AND POWER REQUIREMENTS FOR THE PERIPHERALS HOUSED IN THE COMPUTER MAIN FRAME ARE INCLUDED WITH THOSE FOR THE CPU. A LARGE VARIETY OF STAND-ALONE PERIPHERALS MAY BE ATTACHED TO THE QM-1. THE SPACE AND POWER REQUIRED FOR CERTAIN TYPICAL PERIPHERALS IS LISTED BELOW.

LINE PRINTER - NANODATA P300 (INCLUDING PAPER STAND AND STACKER) 38" DEEP X 30" WIDE X 41" HIGH	120 VAC -	AMPS
CARD READER - NANODATA CR200 (TOP MOUNTED ON STORAGE CABINET) 23" DEEP X 18.5" WIDE X 43" HIGH	120 VAC -	AMPS
CRT - NANODATA DT2 (MOUNTED ON SPECIALLY DESIGNED TABLE) 30" DEEP X 48" WIDE X 43.5" HIGH	120 VAC -	AMPS
TAPE DRIVE - NANODATA TT45 (INDIVIDUAL CABINETS NOT PART OF CPU) 21" DEEP X 26" WIDE X 58" HIGH	120 VAC -	AMPS

CONTACT NANODATA MARKETING DEPARTMENT FOR REQUIREMENTS OF PERIPHERALS NOT SHOWN HERE.

9.6 INSTALLATION LAYOUT -- MEDIUM SCALE CONFIGURATION



CUSTOMERS SHOULD PLAN INSTALLATION DETAILS WITH NANODATA SERVICE PERSONNEL AT LEAST 30 DAYS PRIOR TO DELIVERY. THIS WILL ASSURE PROPER CABLING, CIRCUIT PROTECTION, GROUNDING AND NOISE CONTROL.

APPENDIX A - QM-1 PORT INTERFACE SPECIFICATIONS

A-1 GENERAL SPECIFICATIONS

A-2 THE PORT SIGNALS

A-2.1 GENERAL

A-2.2 FUNCTIONAL DESCRIPTION OF THE PORT SIGNALS

A-2.3 ELECTRICAL SPECIFICATION OF THE PORT SIGNALS

A-2.4 TIMING OF THE PORT SIGNALS

A-2.5 PIN ASSIGNMENT OF THE PORT SIGNALS

A-3 THE INTERFACE CARD

A-1 GENERAL SPECIFICATIONS

THIS SECTION GIVES ELECTRICAL SPECIFICATIONS FOR THE USER DESIRING TO INTERFACE HIS EQUIPMENT DIRECTLY TO THE QM-1 PORT. BECAUSE OF THE UNIQUE ARCHITECTURE OF THE QM-1, IT IS DIFFICULT TO DESIGN EXTERNAL HARDWARE WITHOUT UNDERSTANDING SOME OF THE MACHINE CONCEPTS. THE HARDWARE DESIGNER IS PARTICULARLY ENCOURAGED TO BE FAMILIAR WITH THE FOLLOWING SECTIONS OF THE HARDWARE LEVEL USER'S MANUAL:

- 4.2.5 EXTERNAL STORE
- 4.3 SIX-BIT CONTROL STRUCTURE
- 4.5.2.4 EXTERNAL INTERRUPTS
- 4.6 EXTERNAL INTERFACE
- 5.3 FUNDAMENTAL TIMING CONSIDERATIONS
- 5.5 DATA TRANSFER FUNCTIONS
- 8.1 QM-1 I/O SYSTEM, GENERAL
- 8.2 QM-1 I/O CONTROLS
- 8.3 STANDARD CHANNEL CONTROLLER

THE PORT SIGNALS ARE AVAILABLE THROUGH THE EIGHT PORT CONNECTORS LOCATED AT THE CPU. INTERFACE CARDS THAT MATE WITH THE CPU PORT CONNECTORS ARE AVAILABLE FROM NANODATA. THESE CARDS PROVIDE FOR TERMINATION OF THE USER CABLING TO THE QM-1, AND ALSO HAVE SPACE FOR SOME LOGIC, CABLE DRIVERS, RECEIVERS, ETC. IT IS CONCEIVABLE THAT A SIMPLE DEVICE MIGHT HAVE ALL OF ITS INTERFACE LOGIC MOUNTED ON THE INTERFACE CARD. THIS LIMITS THE CHANNEL TO ONE SUCH DEVICE. THE INTERFACE CARD IS DESCRIBED IN DETAIL IN SECTION A-3. SECTION A-2 GIVES A FULL AND DETAILED DESCRIPTION OF THE PORT SIGNALS.



A-2 THE PORT SIGNALS

A-2.1 GENERAL

SECTION A-2 IS A DETAILED DESCRIPTION OF THE FUNCTIONAL, ELECTRICAL AND TIMING CHARACTERISTICS OF THE PORT SIGNALS. THE PORT SIGNALS ARE SHOWN IN FIGURE 1.

A-2.2 FUNCTIONAL CHARACTERISTICS OF THE PORT SIGNALS

THE FOLLOWING TABLE DESCRIBES THE FUNCTIONAL CHARACTERISTICS OF THE PORT SIGNALS.

NOTES:

- 1) THE LOGICAL LEVEL IS THE RELATIVE VOLTAGE WHEN THE SIGNAL IS "TRUE". A "LOW" SIGNAL, FOR EXAMPLE, IS A SIGNAL WHOSE RELATIVE VOLTAGE IS "LOW" WHEN IT BECOMES "TRUE". (SEE SECTION A-2.3 FOR FURTHER DISCUSSION.)
- 2) CAUTION SHOULD BE TAKEN WITH THE DEFINITIONS OF "PULSE" AND "LEVEL". BECAUSE OF THE UNIQUE ARCHITECTURE OF QM-1 THE PROGRAM HAS FULL CONTROL OVER THE PORT SIGNALS. A "PULSE" IS DEFINED AS A SIGNAL THAT IS TRUE FOR A RELATIVELY SHORT TIME. PROGRAMMING CANNOT AFFECT THE WIDTH OR APPEARANCE OF A "PULSE". A SIGNAL THAT IS DEFINED AS A "LEVEL" CAN BE MADE TO LOOK LIKE A "PULSE", BY ALLOWING THAT LEVEL TO REMAIN CONSTANT FOR ONLY ONE T-PERIOD.
- 3) THE PORT SIGNALS WERE DETERMINED WITH THE NANODATA STANDARD I/O SCHEME IN MIND. THE TERM "NORMALLY USED", THAT OCCURS FREQUENTLY IN THE TABLE BELOW, REFERS TO ITS USE IN NANODATA STANDARD I/O CHANNELS. THESE DESIGNATIONS NEED NOT BE FOLLOWED BY THE USER. THE USER MAY REDEFINE THE FUNCTIONS AND DESIGNATIONS OF PORT SIGNALS TO HIS DEVICE.
- 4) IT IS A NANODATA CONVENTION THAT THE NAMES OF OUTBOUND SIGNALS, THOSE GOING "FROM" THE CPU, ARE PREFIXED BY AN "F" AND INBOUND SIGNALS, THOSE "TO" THE CPU, ARE PREFIXED BY A "T".

SIGNALS FROM THE CPU

SIGNAL NAME	I LOGIC I LEVEL I ----- I	I I SOURCE I ----- I	I I NOTES I
# OF LINES	I SIGNAL I TYPE I	I DESTINATION I	
FDS0-FDS5	I HIGH I ----- I	I CPU G-BUS I ----- I	I SIX DATA LINES FROM THE CPU NORMALLY I USED AS A DEVICE SELECTION CODE TO I ROUTE A COMMAND TO A DEVICE.
6	I LEVEL I ----- I	I ALL PORTS I ----- I	
FFUNC0-FFUNC5	I HIGH I ----- I	I CPU PHANTOM I BUS I ----- I	I SIX DATA LINES FROM THE CPU NORMALLY I USED TO SPECIFY A COMMAND.
6	I LEVEL I ----- I	I ALL PORTS I ----- I	
FXIO	I HIGH I ----- I	I XIO PRIMITIVE I ----- I	I "TRANSMIT I/O" IS A PROGRAMMABLE LEVEL I ITS LENGTH IS DETERMINED BY THE NUMBER I OF CONSECUTIVE T-PERIODS IN WHICH THE I XIO PRIMITIVE IS ACTIVE. IT IS NOR- I MALLY USED TO SELECT THE CHANNEL TO I WHICH DATA OR COMMAND IS SENT.
1	I LEVEL I ----- I	I KA SELECTED I PORT I ----- I	
FXIO STROBE	I HIGH I ----- I	I XIO PRIMITIVE I ----- I	I A STROBE PRODUCED DURING EVERY T-STEP I IN WHICH THE XIO PRIMITIVE IS SPECI- I FIED. NORMALLY USED TO STROBE DATA/ I COMMANDS. ITS TIMING IS SUCH THAT I FDS, FFUNC, AND FDATA ARE VALID DURING I XIO STROBE.
1	I PULSE I ----- I	I ALL PORTS I ----- I	
FRIO	I LOW I ----- I	I RIO PRIMITIVE I ----- I	I "READ I/O" IS USED TO GATE THE TDATA I LINES OF THE KA SELECTED PORT INTO ITS I PORT REGISTER AND NORMALLY USED TO I NOTIFY THE CHANNEL/DEVICE THAT DATA I HAS BEEN READ.
1	I PULSE I ----- I	I KA SELECTED I PORT I ----- I	
FIO CLK	I HIGH I ----- I	I CPU CLOCK I ----- I	I A SYNCHRONIZING PULSE PRODUCED ONCE I EVERY T-STEP INDEPENDENT OF ANY I/O I CONTROLS.
1	I PULSE I ----- I	I ALL PORTS I ----- I	

```

=====I=====I=====I=====
FMC          I LOW      I CPU "MASTER I GENERAL MASTER CLEAR TO ALL QM-1
              I          I CLEAR"      I HARDWARE.
-----I-----I-----I
1            I LEVEL    I ALL PORTS    I
=====I=====I=====I=====
FDATA0 -     I HIGH     I EACH PORT   I THE FDATA LINES ARE THE DIRECT OUTPUTS
FDATA17      I          I REGISTER    I OF THE PORT REGISTERS. THERE ARE
-----I-----I-----I
18           I LEVEL    I EACH PORT   I EIGHT SETS OF FDATA LINES, ONE SET FOR
                                           I EACH SEPARATE PORT.
=====I=====I=====I=====

```

SIGNALS TO THE CPU

SIGNAL NAME	I	LOGIC	I		I
-----	I	LEVEL	I		I
# OF LINES	I	SIGNAL	I	DESTINATION	I
	I	TYPE	I		I

TDATA0 -	I	LOW	I	EACH PORT	I
TDATA17	I		I	REGISTER	I
-----	I	-----	I		I
18	I	LEVEL	I		I
	I		I		I
	I		I		I
=====					
TID0-TID5	I	LOW	I	IO ID AUX	I
-----	I	-----	I		I
6	I	LEVEL	I		I
	I		I		I
	I		I		I
	I		I		I
	I		I		I
	I		I		I
=====					
TATTN DATA	I	LOW	I	ONE OF 30	I
IN	I		I	INTERRUPT	I
	I		I	FLAGS	I
-----	I	-----	I		I
1	I	PULSE	I		I
=====					
TATTN DATA	I	LOW	I	ONE OF 30	I
OUT	I		I	INTERRUPT	I
-----	I	-----	I	FLAGS	I
1	I	PULSE	I		I
=====					
TATTN STATUS	I	LOW	I	ONE OF 30	I
	I		I	INTERRUPT	I
-----	I	-----	I	FLAGS	I
1	I	PULSE	I		I
=====					

A-2.3 ELECTRICAL SPECIFICATIONS OF THE PORT SIGNALS

ALL DRIVERS AND RECEIVERS USED BY NANODATA ARE STANDARD TTL, H OR S SERIES. FIGURE 2 DESCRIBES THE TERMINATION OF SIGNALS AT THE QM-1 PORT. ALL SIGNALS MUST BE BUFFERED AND TERMINATED AS SHOWN BY FIGURE 1. NOT MORE THAN ONE TTL (H SERIES) LOAD (2 MA) MAY BE DRAWN FROM ANY PORT OUTPUT. ELECTRICAL SIGNALS ARE DEFINED AS FOLLOWS:

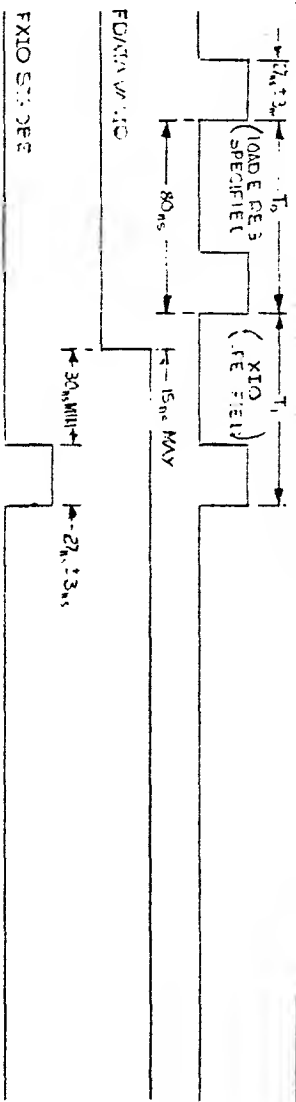
- "HIGH" VOLTAGE - GREATER THAN 2.7 V, LESS THAN 5.5 V.
- "LOW" VOLTAGE - LESS THAN 0.4 V, NOT LESS THAN -1.0 V.

THESE ARE VALUES SUPPLIED BY THE PORT DRIVERS, AND INTERFACE DRIVERS MUST CONFORM TO THEM. A SPACE IS PROVIDED FOR BUFFERING LOGIC ON THE INTERFACE CARD DESCRIBED IN DETAIL IN SECTION A-3.


IMPORTANT NOTE:

RESISTOR TERMINATIONS SHOWN IN FIGURE 2 WITH AN ASTERISK MUST BE INSTALLED ON THE PORT 0 INTERFACE CARD. THEY SHOULD NOT BE USED ON ANY OTHER PORT. IF PORT 0 IS NOT USED, AN INTERFACE CARD WITH TERMINATING RESISTORS MUST BE INSTALLED IN THE PORT 0 CONNECTOR.

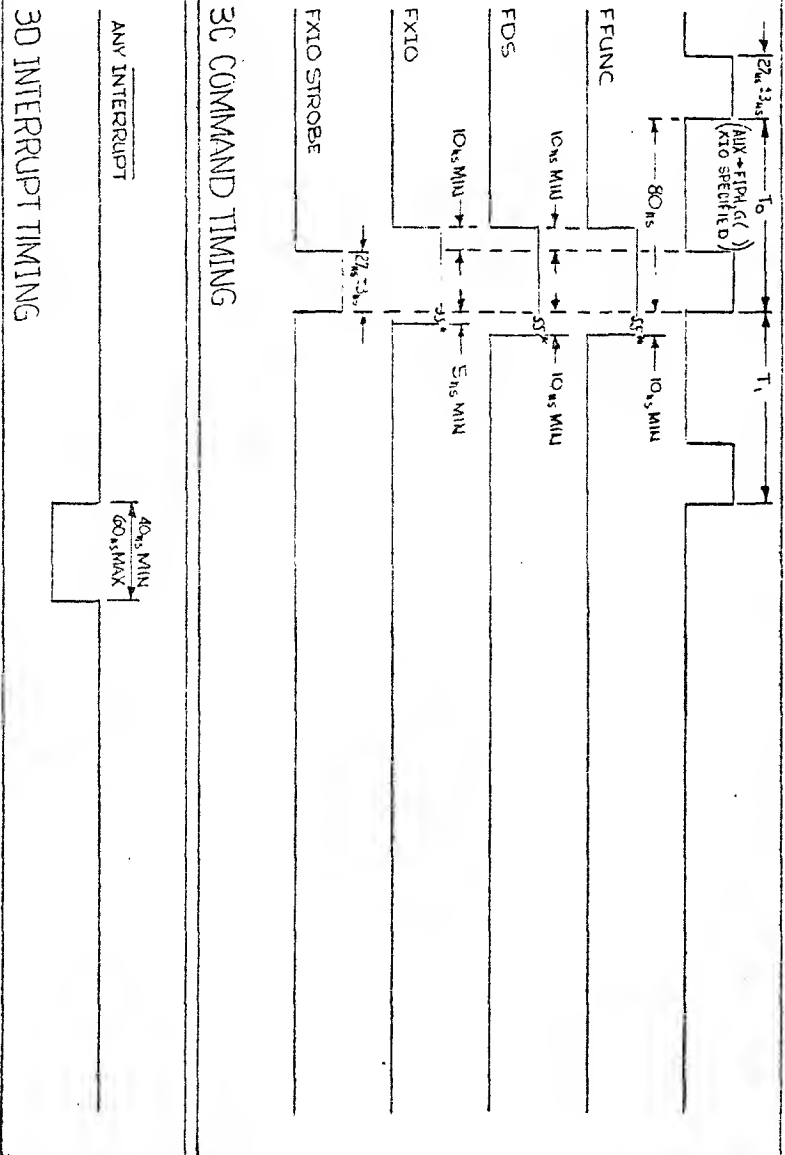
3A DATA OUT TIMING




3B DATA IN & DEVICE ID TIMING



3C COMMAND TIMING



3D INTERRUPT TIMING



NOTE:
 * PROGRAM DEPENDENT
 ** TID HOLD TIME IS REFERENCED TO RIO SINCE NO PORT SIGNAL INDICATES IO.ID → F TRANSFER

FIG 3

A-2.4 TIMING OF THE PORT SIGNALS

FIGURE 3 DESCRIBES THE TIMING OF THE PORT SIGNALS RELATED TO THE OPERATIONS THE SIGNALS ARE INVOLVED WITH.

NOTES:

FIGURE 3A: T1 IS THE EARLIEST T-STEP WHERE XIO CAN BE SPECIFIED IF THE DATA LOADED INTO THE PORT REGISTER IS TO BE VALID.

FIGURE 3B: "INTERRUPT" IN FIGURE 3B MAY BE DATA-IN, DATA-OUT, OR STATUS. INTERRUPT PULSE TIMING IS INDEPENDENT OF CPU TIMING AND MAY OCCUR AT ANY TIME. DATA AND DEVICE ID MUST BE VALID 150 NANOSECONDS FROM THE TRAILING EDGE OF THE INTERRUPT PULSE, SINCE THAT IS THE MINIMUM TIME REQUIRED FOR THE CPU TO RESPOND TO THE INTERRUPT. DATA MUST REMAIN VALID FOR AT LEAST 5 NANOSECONDS AFTER THE TRAILING EDGE OF RIO. SINCE THERE IS NO INDICATION AT THE PORT THAT THE ID IS BEING READ BY THE CPU, ITS HOLD TIME IS REFERENCED TO RIO. NORMALLY, RIO WILL RELEASE THE DATA AND ID.

FIGURE 3C: FXIO STROBE TIMING IS SUCH THAT ALL 3 ELEMENTS OF THE COMMAND (FUNCTION, DEVICE SELECT, AND XIO) CAN BE SPECIFIED IN A SINGLE T-STEP.

FIGURE 3D: AN INTERRUPT OF ANY TYPE CAN OCCUR AT ANY TIME (WITH RELATION TO THE CPU TIMING). THE PULSE WIDTH MUST CONFORM TO THE TIMING LIMITS SHOWN.

A-2.5 PIN ASSIGNMENTS OF THE PORT SIGNALS

TDATA(L)	0 - A3	I	FDATA	0 - A32	I	TID(L)	0 - C6
"	1 - C3	I	"	1 - C32	I	"	1 - A6
"	2 - A10	I	"	2 - A33	I	"	2 - C29
"	3 - C10	I	"	3 - C33	I	"	3 - A29
"	4 - A19	I	"	4 - A35	I	"	4 - C21
"	5 - C19	I	"	5 - C35	I	"	5 - A21
"	6 - A27	I	"	6 - A42	I	FDS	0 - C39
"	7 - C27	I	"	7 - C42	I	"	1 - A39
"	8 - A4	I	"	8 - A44	I	"	2 - C48
"	9 - C4	I	"	9 - C44	I	"	3 - A48
"	10 - A11	I	"	10 - A45	I	"	4 - C47
"	11 - C11	I	"	11 - C45	I	"	5 - A47
"	12 - A20	I	"	12 - A36	I	FFUNC	0 - C49
"	13 - C20	I	"	13 - C36	I	"	1 - A49
"	14 - A28	I	"	14 - A37	I	"	2 - C41
"	15 - C28	I	"	15 - C37	I	"	3 - A41
"	16 - A5	I	"	16 - A38	I	"	4 - C40
"	17 - C5	I	"	17 - C38	I	"	5 - A40

FRIO(L) - A15

FXIO STR - A31

ID SELECT(L) - A16

FIO CLK - A14

TATTN STATUS(L) - A12

FMC(L) - C14

TATTN DATA IN(L) - C12

TATTN DATA OUT(L) - C23

+ 5 V - A9, A26, A43

- 5 V - A7, C7

GND - C1, A18, A34, C52

A-3 THE INTERFACE CARD

THE QM-1 CPU HAS EIGHT, 156 PIN, FEMALE CONNECTORS EACH REPRESENTING A PORT. THE INTERFACE CARD IS A NANODATA STANDARD, UN SERIES, WIREWRAP BOARD WITH A 104 PIN MALE CONNECTOR DESIGNED TO MATE WITH A PORT CONNECTOR IN THE CPU. THIS CARD SIMPLIFIES THE USER'S TASK OF INTERFACING TO THE PORT BY ALLOWING HIM TO PUT HIS DRIVERS, RECEIVERS, AND ANY LOGIC DIRECTLY IN THE QM-1. THE INTERFACE CARD IS AVAILABLE IN THREE DIFFERENT CONFIGURATIONS:

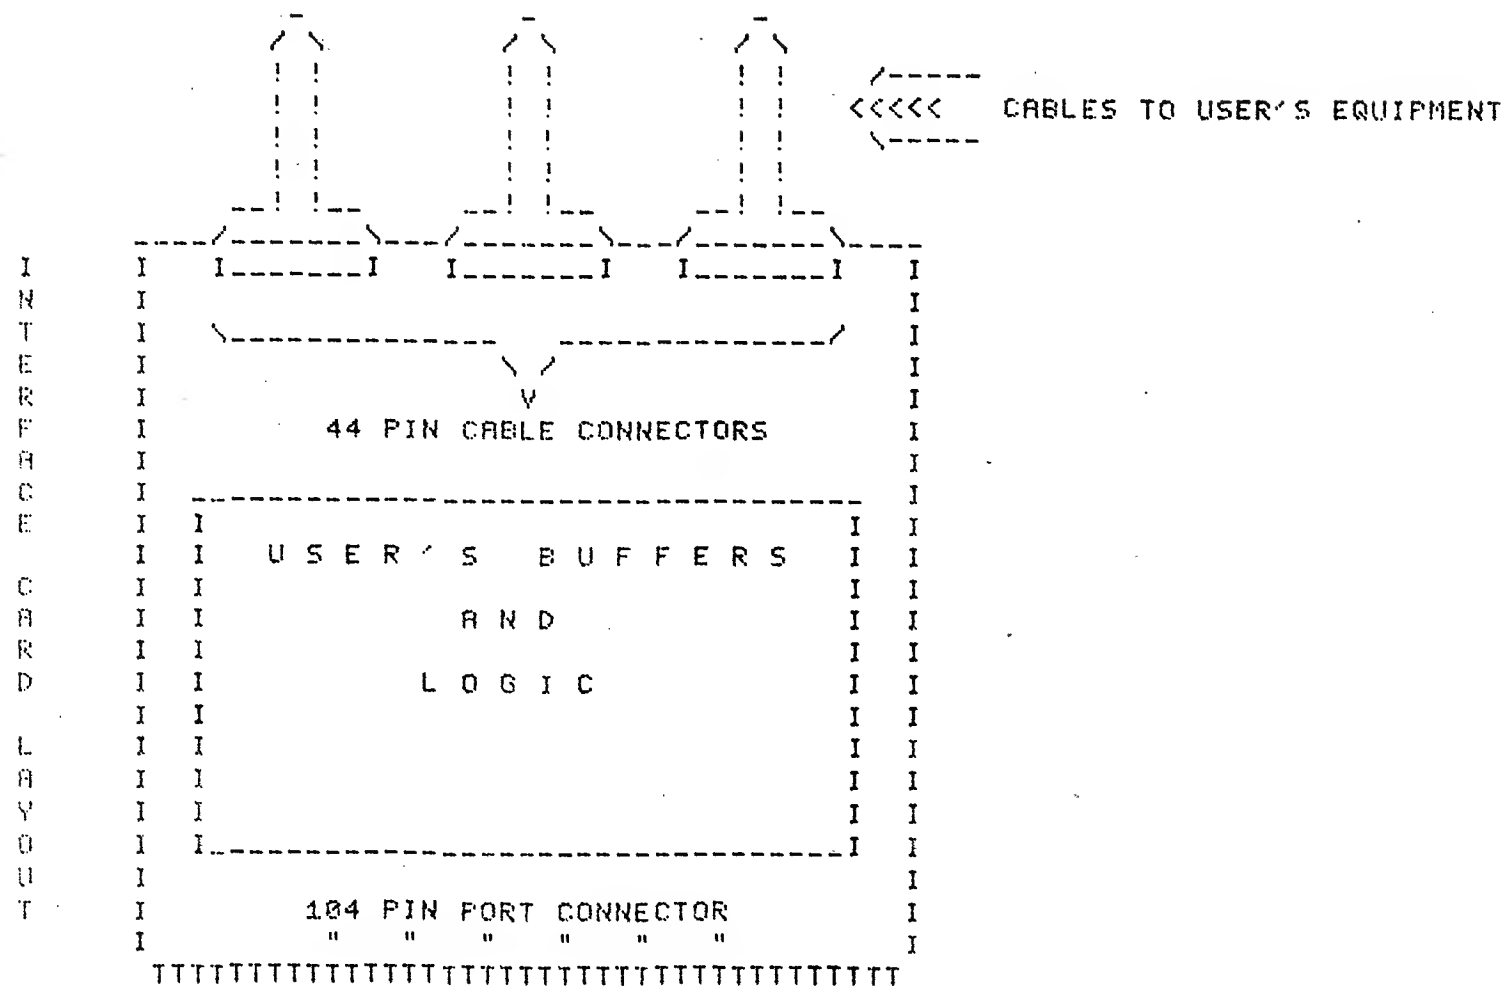
UN 2 - HAS SPACE FOR UP TO 30 14-PIN OR 16-PIN DIPS AND 8 24-PIN DIPS.

UN 3 - HAS SPACE FOR UP TO 12 14-PIN OR 16-PIN DIPS AND 16 24-PIN DIPS.

UN 5 - HAS SPACE FOR UP TO 48 14-PIN OR 16-PIN DIPS.

EACH CONFIGURATION HAS A VOLTAGE PLANE, GROUND PLANE, 104 PIN CONNECTOR, AND ROOM FOR 3 44-PIN CABLE CONNECTORS (CONTINENTAL, MMM 44). +5 VOLT AND -5 VOLT SUPPLIES ARE AVAILABLE AT THE PORT CONNECTOR (SEE SECTION A-2.5). NOT MORE THAN 2 AMPS MAY BE DRAWN FROM THE +5 VOLT SUPPLY AND NOT MORE THAN 400 MA FROM THE -5 VOLT SUPPLY.

IF THE USER DESIRES, NANODATA CAN SUPPLY INTERFACE CARDS COMPLETELY ASSEMBLED, WIRED, AND TESTED TO USER SPECIFICATIONS.



PORT CONNECTOR POSITIONS ON QM-1 BACKPLANE

POSITION #	44	43	42	41	40	39	38	37	
ROW H	-----								/
	I	I	I	I	I	I	I	I	\
	I	I	I	I	I	I	I	I	/
	I	I	I	I	I	I	I	I	\
	I	I	I	I	I	I	I	I	/
	I	I	I	I	I	I	I	I	\
	I	I	I	I	I	I	I	I	/
	I	I	I	I	I	I	I	I	\
	I	I	I	I	I	I	I	I	/
	I	I	I	I	I	I	I	I	\
	I	I	I	I	I	I	I	I	/
	I	I	I	I	I	I	I	I	\
	I	I	I	I	I	I	I	I	/
	I	I	I	I	I	I	I	I	\
	I	I	I	I	I	I	I	I	/
	I	I	I	I	I	I	I	I	\
-----								/	
PORT #	0	1	2	3	4	5	6	7	

APPENDIX B - QM-1 CPU OPTIONAL FEATURES

B-1 CONTROL STORE ADDRESS TRANSLATION AND ACCESS PROTECTION

B-1.1 GENERAL DESCRIPTION

B-1.2 DETAILED DESCRIPTION OF OPERATION

B-1.3 LOADING AUXILIARY MEMORIES

B-1.4 APPLICATIONS

B-1 CONTROL STORE ADDRESS TRANSLATION AND ACCESS PROTECTION

B-1.1 GENERAL DESCRIPTION

THE STANDARD CONTROL STORE OF THE QM-1 IS A FULLY READABLE/WRITABLE 18-BIT WIDE STORE, IMPLEMENTED IN SEMICONDUCTOR MEMORY. IT IS AVAILABLE IN BLOCKS OF 1K WORDS, UP TO A MAXIMUM OF 16K WORDS. CONTROL STORE IS PRIMARILY USED TO HOLD MICROPROGRAMS AND THEIR ASSOCIATED TABLES AND WORK AREAS. IT IS, HOWEVER, A COMPLETELY GENERAL-PURPOSE MEMORY AND MAY BE USED IN ANY WAY APPROPRIATE TO A SPECIFIC APPLICATION. FOR EXAMPLE, CONTROL STORE IS IDEAL FOR USE AS A SCRATCH PAD OR CACHE MEMORY.

IN NORMAL OPERATION, CONTROL STORE IS ADDRESSED FROM ZERO TO THE MAXIMUM INSTALLED MEMORY ADDRESS. ADDRESSES BEYOND THIS RANGE GENERATE ZEROS FOR THE "READ CS" COMMAND, AND CAUSE A NULL OPERATION FOR THE "WRITE CS" COMMAND. ANY EXECUTING MICROPROGRAM HAS ACCESS TO THE FULL RANGE OF INSTALLED ADDRESSES. ALL ADDRESSES MUST BE ABSOLUTE AND NO PORTION OF CONTROL STORE MAY BE EXCLUDED. FOR MANY APPLICATIONS, THIS MODE OF CONTROL STORE OPERATION IS SUFFICIENT. IN PARTICULAR, USERS EXECUTING A SINGLE MICROPROGRAM STREAM WILL PROBABLY HAVE NO NEED FOR A MORE POWERFUL MODE OF CONTROL STORE OPERATION. FOR APPLICATIONS THAT DO REQUIRE A MORE POWERFUL CONTROL STORE FACILITY, THE CONTROL STORE ADDRESS TRANSLATION AND ACCESS PROTECTION OPTION IS AVAILABLE.

WITH THIS OPTION INSTALLED, MICROPROGRAMS HAVE AVAILABLE A VIRTUAL ADDRESS SPACE OF 128K FOR CONTROL STORE. TRANSLATION HARDWARE MAPS THE ADDRESS SUPPLIED BY AN EXECUTING MICROPROGRAM INTO THE ACTUAL ADDRESS SPACE OF 16K (MAXIMUM) THAT CORRESPONDS TO THE CONTROL STORE PHYSICALLY INSTALLED ON THE MACHINE. THE ACTUAL ADDRESS SPACE IS DIVIDED INTO 512 WORD PAGES; 32 SUCH PAGES EXIST IN THE MAXIMUM CONFIGURATION. THE VIRTUAL ADDRESS SPACE IS ALSO DIVIDED INTO 512 WORD PAGES; 256 SUCH PAGES MAY BE REFERED TO BY A MICRO-PROGRAM.

THE TRANSLATION BETWEEN VIRTUAL PAGE AND ACTUAL PAGE IS ACCOMPLISHED BY A SMALL HIGH-SPEED ASSOCIATIVE MEMORY CALLED THE ASSOCIATIVE PAGE SELECTOR. THIS AUXILIARY MEMORY IS LOADED BY A CONTROL PROGRAM WHEN CONTROL STORE IS INITIALLY WRITTEN PRIOR TO RELEASING CONTROL TO THE CURRENTLY EXECUTING MICROPROGRAM. NOW CONTROL STORE ADDRESSES ARE INDEPENDENT OF THE ACTUAL PAGE ADDRESS AT WHICH THE PAGE IS LOADED INTO CONTROL STORE. AN EXECUTING MICROPROGRAM CAN REFERENCE ANY OF ITS CURRENTLY LOADED PAGES. NO CHANGES ARE NECESSARY IN THE NANO-PRIMITIVES USED BY THE MICROPROGRAM.

ANOTHER AUXILIARY MEMORY CALLED THE PAGE ACCESS CONTROL MEMORY IS PROVIDED TO PERMIT CONTROL OVER WHICH PAGES THE CURRENTLY EXECUTING MICROPROGRAM CAN ACCESS. THIS MEMORY IS ALSO LOADED BY A CONTROL PROGRAM PRIOR TO INITIATION OF MICROPROGRAM EXECUTION. THE PAGE ACCESS CONTROL MEMORY CONTAINS A 2 BIT STATUS CODE FOR EACH PHYSICAL PAGE OF CONTROL STORE LOADED AND EACH OF 16 POSSIBLE PARTITION NUMBERS. THE CODE MAY BE SET TO SPECIFY:

- NO ACCESS ALLOWED
- READ ACCESS ALLOWED
- WRITE ACCESS ALLOWED
- FULL ACCESS ALLOWED

IN THIS WAY, CONTROL STORE MAY BE "PARTITIONED" AS APPROPRIATE TO THE TASK NOW BEING EXECUTED. BY A CHANGE IN THE 4 BIT PARTITION NUMBER, THE EXECUTING ENVIRONMENT MAY BE COMPLETELY CHANGED.

WHENEVER CONTROL STORE ACCESS TO THE SPECIFIED PAGE IS RESTRICTED FOR THE CURRENT PARTITION NUMBER, OR WHENEVER THE CONTROL STORE PAGE REFERENCED IS NOT PHYSICALLY PRESENT, AN ADDRESSING EXCEPTION IS GENERATED. THIS PERMITS AN OPERATING SYSTEM TO TAKE THE APPROPRIATE ACTION.

THE NEXT TWO SECTIONS TREAT THE DETAILED OPERATION OF THE CONTROL STORE ADDRESS TRANSLATION AND ACCESS PROTECTION OPTION.

THE 4 BIT PARTITION NUMBER TAKEN FROM THE HIGH ORDER 4 BITS OF FMPC IS INPUT TO A PAGE ACCESS CONTROL MEMORY (PACM) CONSISTING OF 16 BY 64 BITS. THE 64 BITS FOR EACH PARTITION NUMBER CONSIST OF 2 CONTROL BITS FOR EACH POSSIBLE PHYSICAL CONTROL STORE PAGE. THESE 2 BITS ARE CODED TO PERMIT ACCESS TO EACH PAGE AS FOLLOWS:

- 00 - NO ACCESS ALLOWED
- 01 - READ ACCESS ALLOWED
- 10 - WRITE ACCESS ALLOWED
- 11 - BOTH READ AND WRITE ACCESS ALLOWED

AT THE SAME TIME, THE 8 BITS OF THE CONTROL STORE ADDRESS SHOWN IN FIGURE B-1.2A ARE INPUT TO THE ASSOCIATIVE PAGE SELECTOR (APS). THESE 8 BITS SPECIFY THE VIRTUAL PAGE NUMBER OF ONE OF THE POSSIBLE 256 VIRTUAL PAGES. THE 32 BY 8 BIT ASSOCIATIVE MEMORY PROVIDES A SELECTION FROM ONE OF THE 32 POSSIBLE PHYSICAL PAGES OF CONTROL STORE OR INDICATES THAT THE PAGE IS NOT AVAILABLE BY SELECTING NO PAGE.

FINALLY, THE "READ CS" OR "WRITE CS" COMMAND, ALONG WITH THE 64 VALIDATION BITS FROM THE PACM AND THE 32 PAGE SELECT BITS FROM THE APS ARE INPUT TO THE PAGE VALIDATION LOGIC UNIT SHOWN. IF THE VALIDATION BITS SHOW THAT THE ACCESS SPECIFIED IS ALLOWED, THE PAGE SELECT IS SENT TO CONTROL STORE AS THE LOW ORDER 9 BITS OF THE CONTROL STORE ADDRESS SIMULTANEOUSLY SELECT THE WORD WITHIN THE PAGE.

IF THE ACCESS IS NOT VALID OR IF NO PAGE SELECT IS AVAILABLE (INDICATING THAT THE PAGE IS NOT PRESENT IN CONTROL STORE), AN ADDRESSING EXCEPTION IS GENERATED.

B-1.3 LOADING AUXILIARY MEMORIES

BOTH AUXILIARY MEMORIES INVOLVED IN THIS OPTION ARE LOADED BY MEANS OF AN AUXILIARY ACTION COMMAND (SEE SECTION 5.8.2). THIS PERMITS AN EFFECTIVE CONTROL OVER THOSE MICROPROGRAMS THAT SHOULD NOT HAVE THE ABILITY TO MODIFY THE CONTENTS OF THE AUXILIARY MEMORIES.

IN EACH CASE, THE MEMORY ADDRESS IS SPECIFIED BY THE FIPH (PHANTOM) BUS (SEE SECTION 4.6), THE DATA IS PROVIDED AS THE OUTPUT OF THE CONTROL STORE ADDRESS SELECTION MECHANISM (SEE SECTION 5.4.2.1), AND THE APPROPRIATE AUXILIARY ACTION COMMAND IS EXECUTED. THE INTERPRETATION ASSOCIATED WITH EACH MEMORY IS SHOWN BELOW.

ASSOCIATIVE TRANSLATION MEMORY

CONTROL STORE ADDRESS

CELL NUMBER SPECIFIED BY THE FIPH BUS (MODULO 32).
DATA SPECIFIED BY 8 BITS FROM THE CONTROL
STORE ADDRESS AS INDICATED.
AUXILIARY ACTION COMMAND 70.

```
-----
IXI DATA IXXXXXXXXXI
-----
1      8      9
```

BITS

PAGE ACCESS CONTROL MEMORY

CONTROL STORE ADDRESS

CELL NUMBER SPECIFIED BY THE FIPH BUS (MODULO 16),
QUARTER CELL ADDRESS BY HIGH 2 BITS OF CONTROL
STORE ADDRESS AND DATA FROM THE LOW 16 BITS.
AUXILIARY ACTION COMMAND 71.

```
-----
IGCI      DATA      ---I
-----
2          16
```

B-1.4 APPLICATIONS

THE CONTROL STORE ADDRESS TRANSLATION AND ACCESS PROTECTION OPTION IS DESIGNED FOR THOSE SOPHISTICATED APPLICATIONS THAT REQUIRE A MORE POWERFUL MEMORY CAPABILITY AT THE MICROPROGRAMMING LEVEL. THE OPTION:

1. PERMITS SHARING OF RE-ENTRANT MICROCODING.
2. ALLOWS ALL CONCURRENT MICRO-PROCESSES TO HAVE INDEPENDENT ADDRESS SPACES.
3. PROVIDES BASIC CONTROL STORE PROGRAM PROTECTION.
4. PROVIDES CONTROL STORE PARTITIONING BETWEEN 15 INDEPENDENT TASKS.
5. FACILITATES EFFICIENT, HIGH SPEED, SWITCHING BETWEEN RESIDENT MICROPROCESSES.

FOR THOSE APPLICATIONS THAT EITHER REQUIRE, OR CAN MAKE EFFECTIVE USE OF, ONE OR MORE OF THESE CAPABILITIES, THE OPTION SHOULD BE INSTALLED.

COMMENTS REGARDING ERRORS, DEFICIENCIES, OR OMISSIONS IN THIS DOCUMENT WILL BE APPRECIATED. COMMENTS SHOULD BE SENT IN WRITING TO THE TECHNICAL SERVICES MANAGER, NANODATA CORPORATION, 2457 WEHRLE DRIVE, WILLIAMSVILLE, NEW YORK 14221.